



ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY

Methods for Analyzing Electric Load Shape and its Variability

P. Price

Environmental Energy Technologies Division

May 2010



Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

ACKNOWLEDGEMENTS

The work described in this report was coordinated by the Demand Response Research Center and funded by the California Energy Commission (Energy Commission), Public Interest Energy Research (PIER) Program, under Work for Others Contract No. 500-03-026 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The author thanks Mary Ann Piette, Sila Kiliccote, Johanna Mathieu, and Jessica Granderson at Lawrence Berkeley National Laboratory for unusually helpful direction and discussion.

ABSTRACT

Current methods of summarizing and analyzing electric load shape are discussed briefly and compared. Simple rules of thumb for graphical display of load shapes are suggested. We propose a set of parameters that quantitatively describe the load shape in many buildings. Using the example of a linear regression model to predict load shape from time and temperature, we show how quantities such as the load's sensitivity to outdoor temperature, and the effectiveness of demand response (DR), can be quantified. Examples are presented using real building data.

Keywords: building energy consumption, whole-building electricity, electric load shape

Please use the following citation for this report:

Price, Phillip N. (Lawrence Berkeley National Laboratory). 2010. *Methods for Analyzing Electric Load Shape and its Variability*. California Energy Commission. Publication number: CEC-XXX-2010-XXX.

TABLE OF CONTENTS

Acknowledgements	i
ABSTRACT.....	ii
EXECUTIVE SUMMARY	1
Introduction	1
Graphical approaches.....	1
Limitations of Graphical Displays	2
Parameters that Describe Load Shape.....	3
Typical office building load shape in summer	3
Summary Statistics Without Weather Adjustment	3
Adjusting for Weather and Other Factors	4
Summarizing errors in predicted electric load	5
Linear regression models for load shape.....	7
Estimates of Demand Response effectiveness	11
Conclusions	12
CHAPTER 1: Graphical approaches to identifying changes and anomalies in electric load.....	15
Limitations of Graphical Displays	18
CHAPTER 2: Parameters that Describe Load Shape	20
Typical office building load shape in summer	20
Summary Statistics Without Weather Adjustment	21
Adjusting for Weather and Other Factors	25
Summarizing errors in predicted electric load	28
CHAPTER 3: Linear regression models for load shape.....	30
Estimates of Demand Response effectiveness	41
Quantifying changes with time.....	44
CHAPTER 4: Conclusions.....	53
References.....	54

LIST OF FIGURES

Figure 1: Electric load versus time (at 15-minute intervals) for four buildings in Martinez, California, for June 2008.....	2
Figure 2: Stylized load shape for a typical office building in summer.....	3
Figure 3: Sketch indicating five parameters that we recommend as a minimum set to characterize load shape.....	4
Table 1: Measures of model fit for linear regression prediction of daily peak load in Building A and Building C, in summer 2008. The model includes only day of week and outdoor temperature.....	6
Figure 4: Predictions (thick line) and observations (thin line) of daily peak electric load for several months beginning May 1, 2008.....	7
Figure 5: Predicted (thick line) and measured (thin line) electric load in Juvenile Hall for four days in June, 2008. Data and predictions are for 15-minute intervals. Residuals, i.e. the difference between observed and predicted load, are also shown.....	8
Figure 6: Load predictions and measurements for Office Building A, for 15-minute intervals, for several days in June 2008.....	9
Figure 7: Stylized plot of expected thermal conditioning load vs. outdoor temperature	10
Figure 8: Excess load vs. temperature, for an office building in Martinez, California. Estimates are from a model that includes a piecewise-linear relationship between load and temperature.....	11
Figure 9: Predicted (thick line) and observed (thin line) load in an office building on six DR days in summer 2008. In each case, the prediction ignores the effect of DR. The vertical dotted lines at 12:00 and 18:00 indicate the DR period.....	12
Figure 10: Electric load versus time (at 15-minute intervals) for four buildings in Martinez, California, for June 2008.....	15
Figure 11: Electric load for June 1-7 inclusive (solid line) and June 22-29 (dashed line), in Juvenile Hall (top) and county jail (bottom).	16
Figure 12: Load profile by day of week, averaged over the entire summer of 2008 (thick line) and load from the first week in June (thin line), for four buildings.....	18
Figure 13: Stylized load shape for a typical office building in summer.....	20
Figure 14: Actual load versus time for an office building in Martinez, California in July 2008...	21
Figure 15: Sketch indicating five parameters that we recommend as a minimum set to characterize load shape.....	22
Table 2: Parameters that characterize electric load shape (for a single day) in a typical office or retail building.....	23
Figure 16: Nighttime near-base load (lower curve) and daytime near-peak load (upper curve) by day in the same four buildings whose load is shown in Figure 10 for the spring, summer, and fall of 2008.....	24
Table 3: Median daily near-base load, median daily near-peak load, Coefficient Of Variation (COV) of daily peak load, and median COV of load during the high-load period (i.e. the	

period when load is closer to the peak than to the base), for four buildings in Martinez, California, for days in spring-fall 2008.....	25
Figure 17: Predictions (thick line) and observations (thin line) of daily peak electric load for several months beginning May 1, 2008.....	28
Table 4: Measures of model fit for linear regression prediction of daily peak load in Building A and Building C, in summer 2008. The model includes only day of week and outdoor temperature.....	29
Figure 18: Temperature-adjusted baseline for Juvenile Hall, for each 15-minute interval throughout the week.	31
Figure 19: Temperature-adjusted baseline, and the load predicted by the model described in Equation 5, for Juvenile Hall, Sunday-Wednesday of the second week of June 2008	32
Figure 20: Predicted (thick line) and measured (thin line) electric load in Juvenile Hall for four days in June 2008. Data and predictions are for 15-minute intervals. Residuals, i.e. the difference between observed and predicted load, are also shown.	33
Figure 21: Predicted and observed hourly average load in Office Building A, for the same days as shown in the previous figure. Residuals are also shown.	33
Figure 22: Load predictions and measurements for Office Building A, for 15-minute intervals, for several days in June 2008.	36
Figure 23: Load predictions and measurements for Office Building A, for 15-minute intervals, for several days in June 2008..	36
Figure 24: Stylized plot of expected thermal conditioning load vs. outdoor temperature.	38
Figure 25: Stylized plot of expected thermal conditioning load vs. outdoor temperature, showing the effect of decreasing chiller efficiency as the chillers approach full capacity. ...	38
Figure 26: Excess load vs. temperature, for an office building in Martinez, California. Estimates are from a piecewise-linear model summarized in Equation 9.....	41
Figure 27: Predicted (thick line) and observed (thin line) load in an office building on six DR days in summer 2008. In each case, the prediction ignores the effect of DR.....	42
Figure 28: Predicted (thick line) and observed (thin line) load by hour of day in a manufacturing and office building, on eight days in 2008 during which Demand Response was performed from 12:00 to 18:00.....	43
Figure 29: Predicted (thick line) and observed (thin line) load by hour of day in the same building that is considered in Figure 28 for a week in 2008 that included three DR days....	43
Figure 30: Unexplained weekly load shift in the County Jail during the summer of 2008, from the model summarized in Equation 9.	44
Figure 31: Electric load at 15-minute intervals for an office building in California, for the first full work week in June 2007 (top) and June 2008 (bottom). Monday is at the far left.	45
Figure 32: High-load duration in the office building of Figure 32Error! Reference source not found., for weekdays in 2007 and in most of 2008..	46
Figure 33: Daily electric energy (kWh) at Ikea Palo Alto for each day from April 30, 2007 through December 31, 2009..	47

Figure 34: Daily maximum load and minimum load. A horizontal line at 800 kW helps compare load across years.....	48
Figure 35: Hours of high load, by day of the year, for 2007 (top), 2008 (middle) and 2009 (bottom).	49
Figure 36: Electric load vs time for two-week periods beginning with the first Sunday in June 2007, 2008, and 2009. Dashed lines at 400 and 1000 kW facilitate comparing one year to another.	50
Figure 37: Electric load for the first two full weeks of June 2008 (thick line) and 2009 (thin line). Load is lower in 2009 than in 2008, but how much of this is weather-related?	51
Figure 38: Model predictions for two weeks in June (starting with the first Sunday) for 2008 (thick line) and 2009 (thin line).	52

LIST OF TABLES

Table 1: Measures of model fit for linear regression prediction of daily peak load in Building A and Building C, in summer 2008.....	6
Table 2: Parameters that characterize electric load shape (for a single day) in a typical office or retail building.	23
Table 3: Median daily near-base load, median daily near-peak load, Coefficient Of Variation (COV) of daily peak load, and median COV of load during the high-load period (i.e. the period when load is closer to the peak than to the base), for four buildings in Martinez, California, for days in spring-fall 2008.....	25
Table 4: Measures of model fit for linear regression prediction of daily peak load in Building A and Building C, in summer 2008.....	29

EXECUTIVE SUMMARY

Introduction

“Whole-building electric load” is the total electrical power used by a building at a given moment. The load changes with time in response to changes in lighting levels; heating, ventilating, and air conditioning (HVAC) requirements; and uses such as computers, copy machines, and so on. The curve that represents load as a function of time, called the “load shape,” can often yield useful information. Unexpectedly high night-time loads may indicate waste (such as lights that needlessly remain on when the building is unoccupied); a change in load shape may indicate an equipment or thermostat malfunction; unexpectedly high sensitivity to outdoor temperature may indicate that excessive outdoor air is being brought into the building by the HVAC system; and so on.

In this report, we discuss several elements of electric load shape analysis:

1. Characterizing daily load shape: what is a small set of parameters that are useful for describing the load variation during a day, and from one day to the next?
2. Describing energy consumption changes over long timescales (months or years): has the energy consumption changed? If so, was the change gradual or sudden?
3. Relating changes in energy consumption to explanatory variables. To what extent is higher energy use associated with higher outdoor temperatures? Did consumption increase at night or during the day? On weekdays or weekends? Was demand response effective?

We begin by making a few suggestions concerning graphical displays of load data. We then define some terminology to describe load shapes, and introduce several ways of describing load shapes statistically, with examples from real data. Weather sensitivity is then discussed, along with several standard approaches to adjusting for weather in load predictions. We choose linear regression modeling to illustrate weather adjustments, begin with simple temperature standardization and moving on to more sophisticated approaches. Methods for quantifying demand response effectiveness are also discussed. Finally, we give several examples to illustrate how the methods in this paper can be used to detect and quantify changes in building behavior.

Graphical approaches

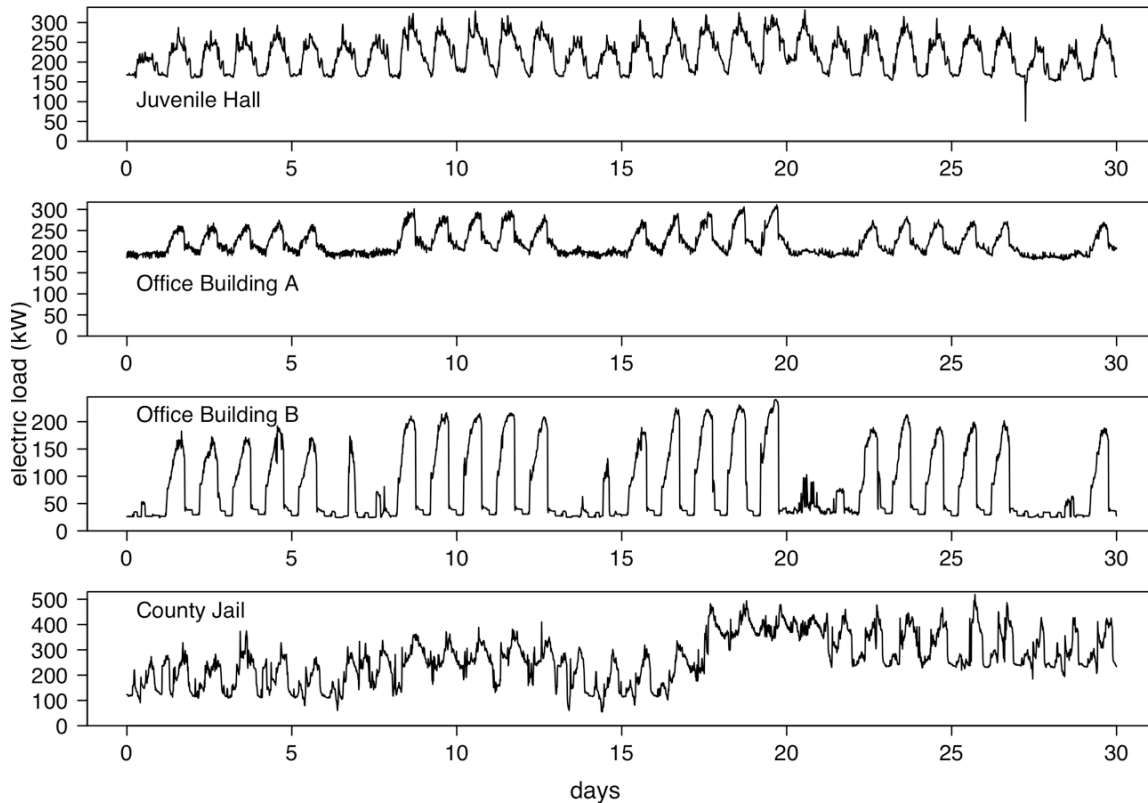
There are some helpful general principles when overlaying load shape curves:

1. Displayed in black and white, two or three curves are often the most that can be shown without becoming visually confusing, but with the use of color this can be increased to four or five.
2. The y-axis should always start at zero.
3. Plots that show a time period of a few days, up to about a week, are best for comparing one time period to another. Longer periods require the plot to be too compressed along the time axis.
4. Choose the plot aspect ratio so that major features have a slope of between 30 and 60 degrees up or down, as is the case for Figure 12. Figure 1 violates this principle slightly: the morning rise in load and the evening fall in load are too steep. As Cleveland (1994) has noted, people are best at comparing slopes near 45 degrees, so very steep or shallow slopes can make it hard to notice asymmetries between the rate of rise and the rate of fall, and to make other useful comparisons. For Figure 1, the solution would be to show

fewer days of data on each plot, but in the present paper we prefer to present more data and are willing to violate the rule of thumb concerning the slopes.

Sometimes one or more of the principles listed above must be violated, in order to encompass a time period of interest or for other reasons. Figure 1, which shows electric load in four different buildings for the same month, shows more days than is ideal, resulting in load curves that are too steep (violating points 3 and 4).

Figure 1: Electric load versus time (at 15-minute intervals) for four buildings in Martinez, California, for June 2008.



Limitations of Graphical Displays

Although simply overlaying load data from different time periods can be useful for noticing and characterizing changes, there are many limitations to this approach. One problem is that weather-related variation from one period to the next is expected – specifically, energy use is expected to be positively correlated with outdoor temperature (and humidity) during summer, and to be negatively correlated with outdoor temperature during winter.

Another limitation of relying on graphical comparison of load shape in different periods is that some phenomena may be difficult to recognize. For instance, noticing a tendency for load to increase gradually over a long period might be quite difficult, since this small trend will often be superimposed on seasonal variation and other features.

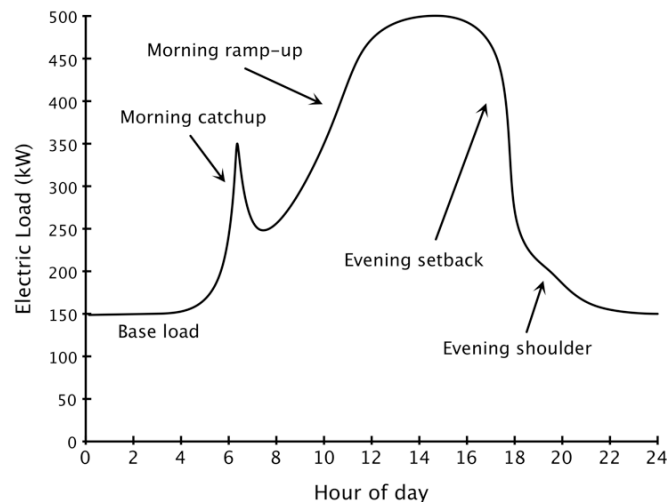
Finally, and most importantly, graphical approaches to comparing load shapes are useful only when people are able to devote time to using them. Automated analysis methods can potentially reduce the effort required to identify problems or opportunities for energy savings.

Parameters that Describe Load Shape

Typical office building load shape in summer

As the examples above illustrate, building load shapes can vary widely. However, many office building load shapes share some characteristics. Figure 2 shows a stylized summer office load shape.

Figure 2: Stylized load shape for a typical office building in summer.

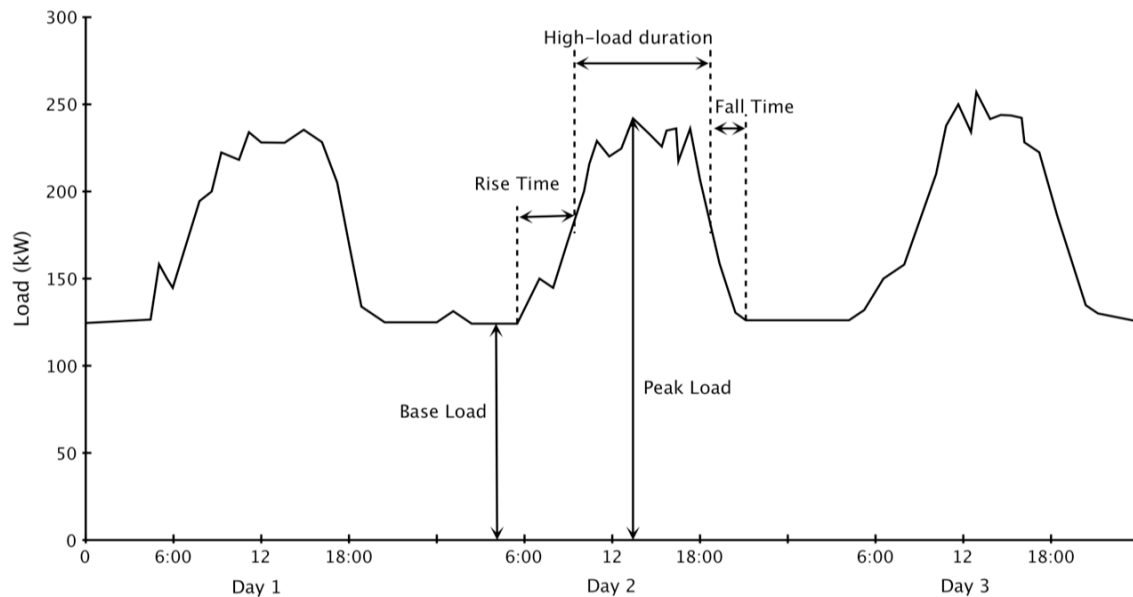


Summary Statistics Without Weather Adjustment

As with any time series, load curves invite the calculation of a wide variety of summary statistics. It is useful to distinguish between two kinds of summary statistics: those that summarize various aspects of the load and its variability, and those that summarize the amount of load or load variability that is not related to temperature or humidity. For instance, consider a building's base load. The base load might vary from month to month, but this variation might be almost perfectly predicted by the nighttime outdoor temperature. For some purposes, a data analyst might care only about the variability, whatever the cause; for others, they might be interested only in the amount of variability that is not related to weather.

Figure 3 shows several parameters that can be used to summarize load shape. Each day, the base load and peak load are calculated, as well as three time intervals. The first time interval quantifies how long it takes for the load to increase halfway to the peak load (i.e. to reach a load closer to the maximum load than to the base load); we call this the "rise time." The second quantifies how long the load stays above the halfway mark; we call this the "high-load duration." And the third, the "fall time," quantifies how long it takes the load to fall from the halfway point back to the base load. Additionally, we summarize the variability of the load during the high-load period, by calculating the standard deviation of the load during that time.

Figure 3: Sketch indicating five parameters that we recommend as a minimum set to characterize load shape. The value of each parameter can be calculated for each day; these values can then be summarized (e.g. average peak load, standard deviation of peak load).



Some buildings, or at least building data, have a peak load – the very highest 15-minute data point – that is substantially higher than any other data point. In some cases we are interested in that highest data point, but more often we are interested in the behavior of the building at slightly longer timescales. For the latter purpose, we define a daily “near-peak load” to be the 97.5th percentile load, i.e. a value that is greater than 97.5% of the load measurements that day. In many buildings and for many purposes, this value is both more stable (i.e. exhibits less day-to-day variability) and more relevant than the peak value. Although we recommend using the near-peak load rather than the true peak, some building operators (or researchers) may be more interested in the actual peak, because some pricing programs have charges or penalties based on the absolute daily peak. (In general we recommend against using the absolute peak in a pricing program. To see why, consider two buildings that both have 15-minute peak load periods during which they use exactly the same amount of electricity. In one of these buildings, the load is entirely during a single 15-minute billing period; in the other it is split across billing periods, with 7.5 minutes in each. These buildings may have substantially different values for their 15-minute peak load, and thus may face different pricing levels, even though from a utility perspective there is little reason to charge one more than the other).

The schematic load shape shown in Figure 2 leaves little doubt as to when the load begins to increase from the base load, or returns to the base load; this is the case for many real buildings as well. However, in some buildings the situation is more complicated (see data from the county jail in Figure 1, for example). Indeed, some buildings, such as industrial facilities, may have load shapes that look completely unlike the schematic load shape in Figure 3 and may require different descriptive parameters.

Adjusting for Weather and Other Factors

There are at least two types of load prediction: (1) *load forecasts* are predictions of future energy use, and may be used for advanced building control, demand-response (DR) planning, and other purposes; and (2) *post-facto predictions* can be made for comparison with historical data to identify anomalies and faults, to determine the effect of changes in building operation, to verify retrofit effectiveness, to determine the temperature sensitivity of a building’s energy

consumption, for DR evaluation, or for other purposes. Load forecasting is more challenging than post facto prediction, because forecasting is necessarily based entirely on data that are available prior to the time period being forecast. For instance, consider predicting peak energy consumption for the first day of the spring in which cooling is needed: the forecast must be made without any recent data that include cooling load, whereas a post-facto analysis can include data collected later in the spring. In this report, we focus primarily on post-facto analysis rather than on forecasts.

Several methods have been applied by researchers or as part of Energy Information Systems (EIS) to determine mathematical relationships that predict load. Granderson et al. (2009) noted that the following methods are in current use: (1) Linear regression of load on outdoor temperature and humidity, with indicator variables for time of day or time of week; (2) Non-linear models such as locally weighted regression; (3) neural network models; (4) “bin” models in which conditions (such as outdoor temperature) are divided into discrete bins, and load predictions are based on the average load for time periods that share the same bin as current conditions; (5) nearest-neighbor methods, in which the current load is predicted to be the same as it was when previous conditions were closest to current conditions. Any of these prediction methods can be used for several purposes, including fault detection (by finding times when the building isn’t behaving as it has behaved in the past) and quantifying the effectiveness of Demand Response (DR). In each case, the predicted load is compared to the actual load.

In this report, we discuss various linear regression models because – when constructed appropriately – they provide a good fit to load data in most buildings; their results are easy to interpret; they are easy to modify; and they present modest computational burden. This latter point is important, even with today’s computers, when working with load data collected at 15-minute intervals: one year of data comprises about 35,000 data points for each variable in the model, which makes the full analysis large enough that some software products can struggle with it or fail altogether.

Consider the model described by **Equation 2**, where \hat{y} is the predicted peak load for a given day of the week and a given temperature:

$$\hat{y}(\text{day of week}, T(t)) = \beta_{\text{day of week}} + \beta_T T(t) \quad \text{Equation 1}$$

where the β values are regression coefficients. We make the time-dependence of temperature explicit for convenience in future equations. In this model, each day of the week has its own typical peak load, and the load increases proportionally with temperature.

For three of the buildings featured in Figure 1 this model works very well for predicting summer daily peak load, but for the jail the fit is rather poor. Figure 4 shows peak load for each day in summer, as predicted (thick line) and measured (thin line) for the four buildings.

Summarizing errors in predicted electric load

A standard way of summarizing the degree of fit between predictions and data is to quote the value of r (the correlation) or R^2 (the squared correlation, which is the ratio of variance in the data that is explained by the model to total variance in the data). For either of these, a value near 1 means good agreement between prediction and data, and a value near 0 means poor agreement between prediction and data. However, these are probably not particularly relevant summary statistics for evaluating the usefulness of models for DR. The problem is that the r - and R^2 values depend on the variance of the data, including (in this case) the weekday and weekend variability. This is illustrated by comparing Juvenile Hall and Building B: Building B is much more variable, because of its large “weekend effect” (load is much lower on weekends than on weekdays). This large amount of variability is captured by the model, since it includes a different base load for each day of the week, so the model explains most of the variance in

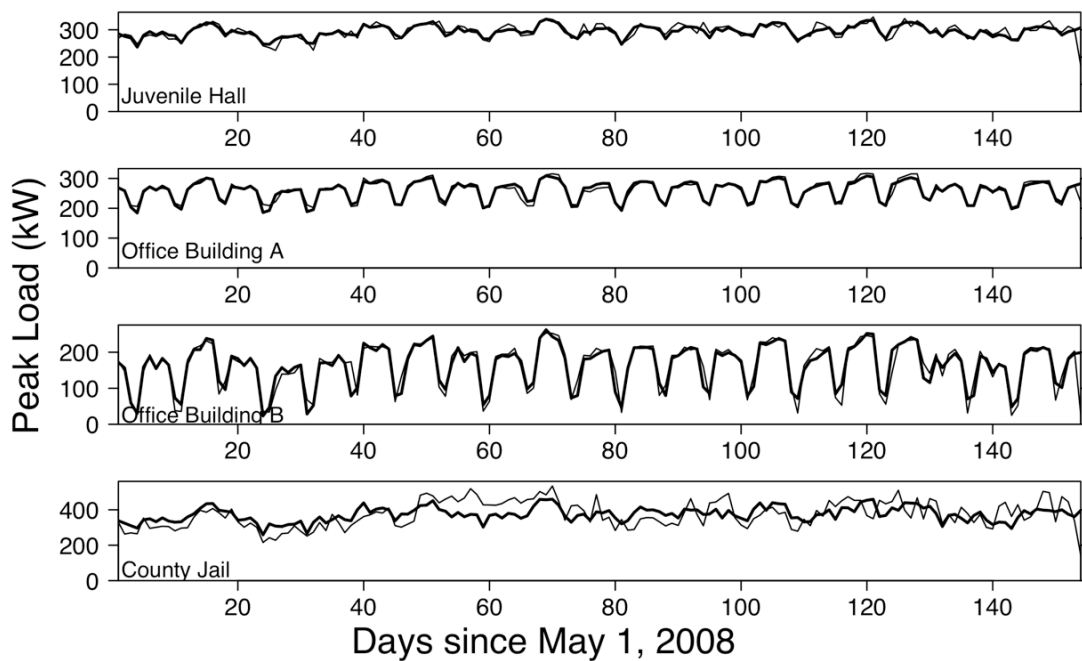
Building B's load. In contrast, Juvenile Hall is much less variable from day to day, and random (unexplained) variability is a larger fraction of the total variability than is the case in Building B. As a result, R^2 is higher in Building B than in Juvenile Hall, as shown in Table 1. The other statistics in the paper, Standard Error and Geometric Mean of Relative Absolute Error (GMRAE), are discussed in the body of this report.

Table 1: Measures of model fit for linear regression prediction of daily peak load in Building A and Building C, in summer 2008. The model includes only day of week and outdoor temperature.

Building	Standard error (kW)	R^2	GMRAE (%)
Juv. Hall	15.3	0.68	2
A	9.9	0.92	1
B	21.9	0.88	5
Jail	62.6	0.31	11

For DR purposes, summaries of model fit other than r or R^2 may be more useful, either for comparing how different models behave or for comparing how a single model behaves across different buildings. The standard error is one simple metric that can be useful. In Table 1 we see that the linear model usually predicts the peak load in Juvenile Hall within about 15 kW, and predicts the peak load in Building B within about 22 kW. In these buildings, which are at least roughly comparable in terms of typical power consumption, simply comparing the standard errors makes sense. But to compare model fit for buildings that differ widely in load, it is useful to have a metric that takes the typical load into account.

Figure 4: Predictions (thick line) and observations (thin line) of daily peak electric load for several months beginning May 1, 2008.



Linear regression models for load shape

Divide a week into intervals (indexed by i); for instance, if electric load data are available for 15-minute intervals, the first interval is from midnight to 12:15 on Sunday morning, the second interval is from 12:15 to 12:30, and so on. A different regression coefficient for each period allows each period to have a different predicted load. Additionally, we expect that for all intervals there is a temperature effect: when the outdoor temperature is high, cooling load will increase with temperature, and when the outdoor temperature is low, heating load will increase as temperature decreases. For some range of moderate temperatures, the load may be insensitive to temperature because neither cooling nor heating is needed. The outdoor temperature at which cooling (or heating) is needed can be estimated if there is enough data from a building, although this estimation adds complexity. For example, load may increase as outdoor temperature increases when the temperature is over 65 F, and may increase as outdoor temperature decreases when the temperature is under 55 F.

Figure 5: Predicted (thick line) and measured (thin line) electric load in Juvenile Hall for four days in June, 2008. Data and predictions are for 15-minute intervals. Residuals, i.e. the difference between observed and predicted load, are also shown.

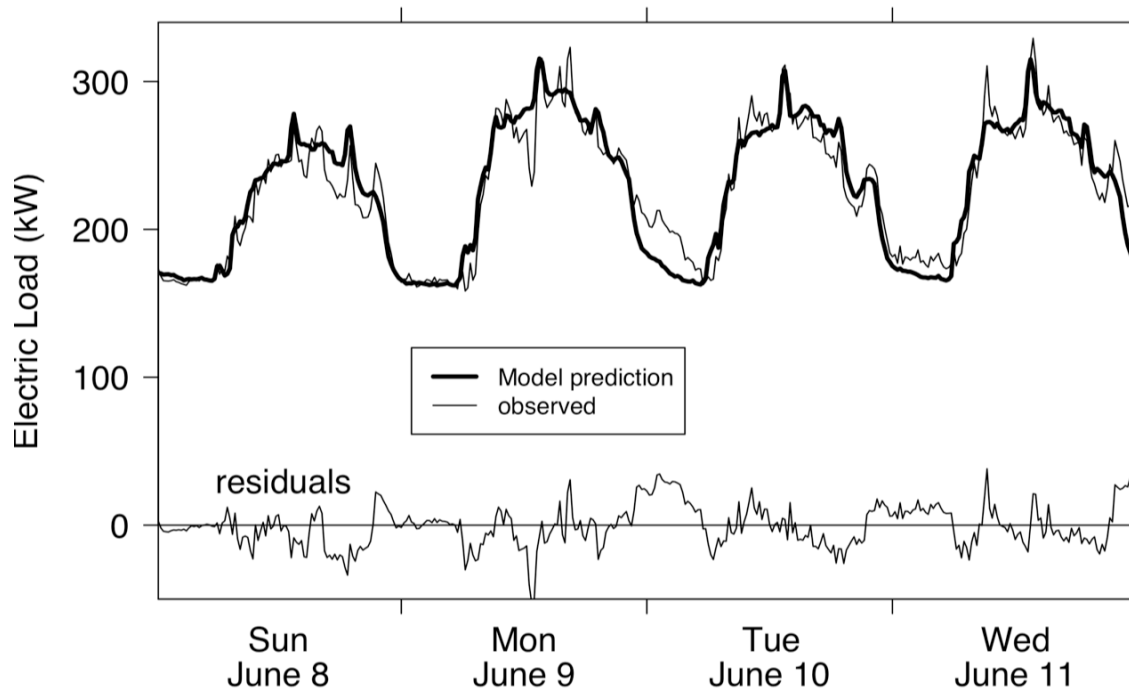
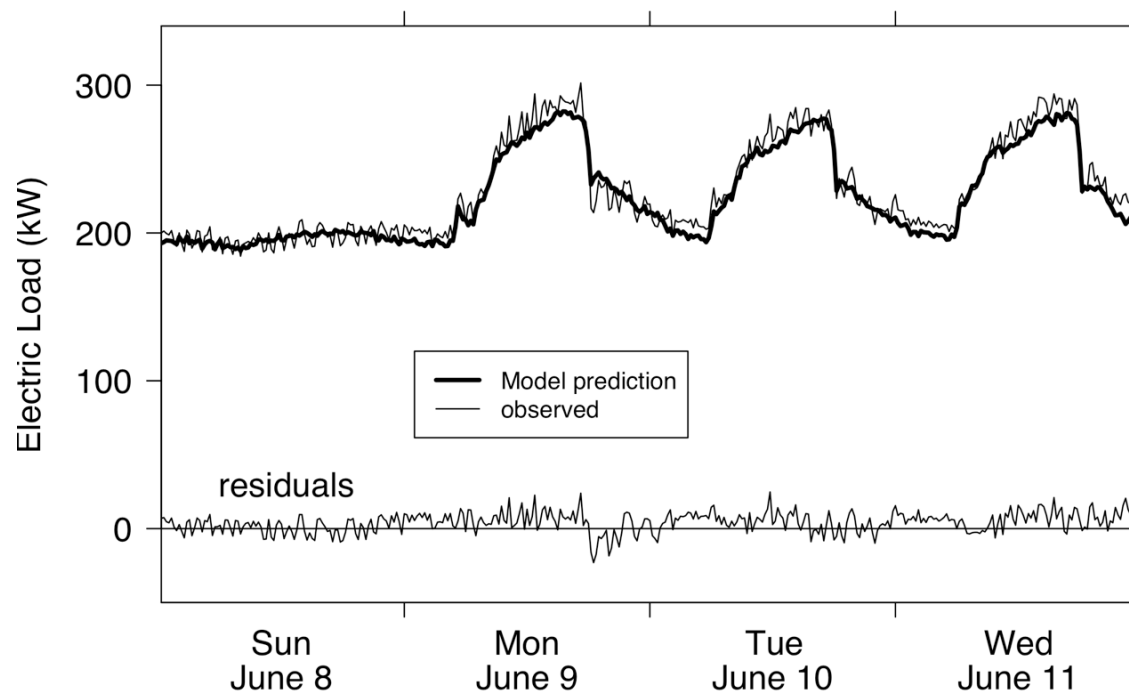


Figure 5 shows the model predictions along with the actual load data for Juvenile Hall from a portion of the second week of June 2008. The difference between measurements and predictions, as a function of time, is shown at the bottom of the figure. The data (thin line) show a prominent peak at 1:30 PM each day, which indeed occurs almost every day at that time. But the data also show sharp peaks and valleys throughout the day that are not predicted from the simple model.

As the above example shows, even a simple regression model that incorporates time of week and outdoor temperature performs quite well for Juvenile Hall. However, there is one obvious model change that might be necessary to improve the model performance for buildings that (unlike Juvenile Hall) are unoccupied or sparsely occupied for part of each week. For such buildings, it makes sense to allow the temperature sensitivity (the coefficient β_{Tover}) to differ between nighttime and daytime. This should improve the fit because, if the temperature setpoint changes during the day --- if the indoor temperature is allowed to be higher at night than during the day --- the temperature sensitivity should vary as well. To see why this is so, consider a case in which the interior of the building is held at 70F during the day, and is allowed to climb to 78F at night and on weekends. An outdoor temperature of, say, 72 F might lead to substantial daytime cooling requirements because of both the need to cool the ventilation air and because of internal loads, but the same temperature might lead to no cooling requirement at all at non-working hours (such as weekends or nighttimes) because the internal loads are lower and because the setpoint temperature is higher.

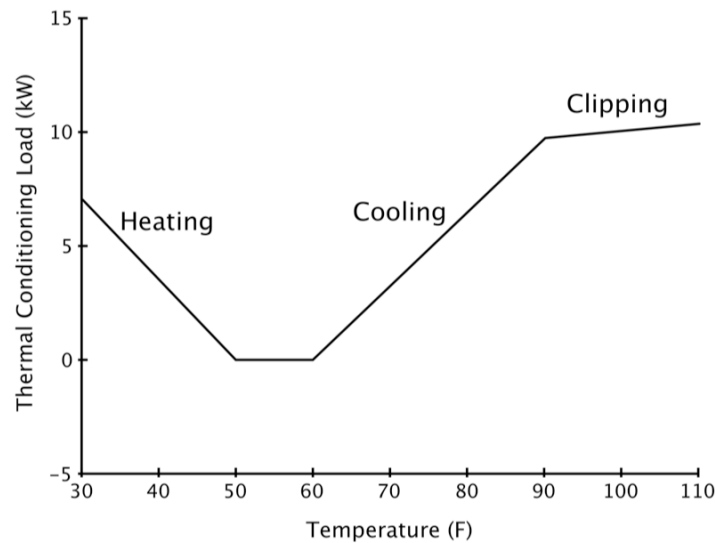
Figure 6 compares observations to predictions from a model that includes time of week and also two temperature sensitivity coefficients, one for evenings and weekends and one for workdays. The regression coefficients indicate that during working hours, the building uses an additional 1.34 kW for each degree F that the outdoor temperature exceeds 65 F, whereas at other times it uses only 0.33 additional kW per degree F.

Figure 6: Load predictions and measurements for Office Building A, for 15-minute intervals, for several days in June 2008. Predictions are made using a model that estimates outdoor temperature sensitivity separately for time intervals when the building is and is not occupied.



Simply allowing the effect of temperature on load to differ between working hours and other periods does not capture all of the systematic variation in temperature sensitivity that is expected. In fact, even at a given time of day the expected temperature dependence may be rather complicated. Figure 7 is a stylized plot of the expected relationship between load and outdoor temperature in a hypothetical building; at any given point, the slope of the relationship is what we call “temperature sensitivity.” At low temperatures, the lower the outdoor temperature, the more energy the heating system consumes (e.g. for an electric heat pump); in this regime, electric load decreases with temperature. At cool temperatures, the building needs neither cooling nor heating. For moderate outdoor temperatures, the cooling load increases as the temperature increases, due to chiller operation, fan power, etc. In extremely hot conditions the cooling plant may be undersized, in which case its energy consumption will level out: once the system is working as hard as it can, its energy consumption cannot increase, a phenomenon we call “maxing out” or “clipping.” (However, there could be additional loads, e.g. from the use of desk fans, that continue to increase with temperature).

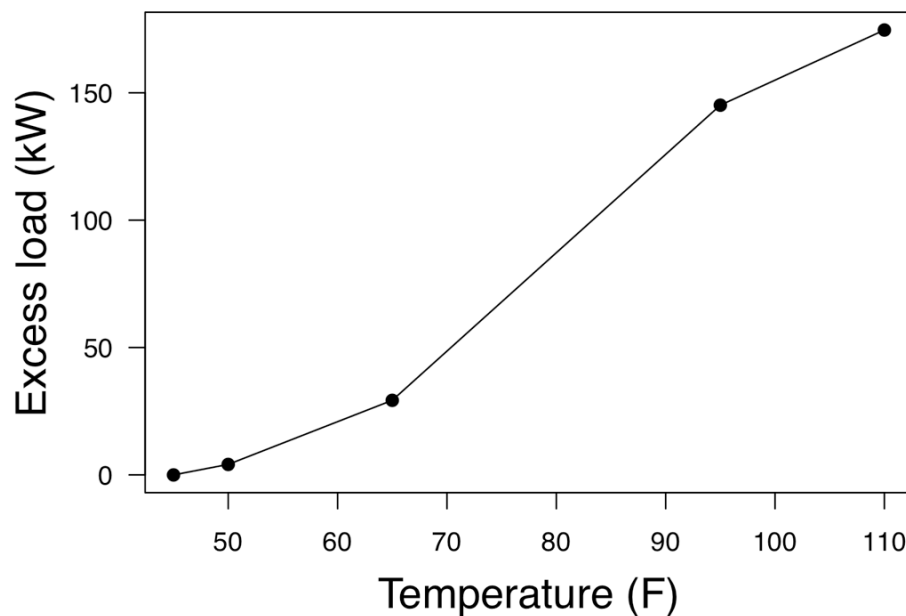
Figure 7: Stylized plot of expected thermal conditioning load vs. outdoor temperature



The effects of maxing out and varying chiller efficiency are both minor from the standpoint of predicting long-term (weekly or monthly) building energy consumption: the effects are important only when the cooling system is operating close peak capacity, which is likely to happen only a few times per year because most cooling systems are sized to meet the demand on the hottest days.

However, although small on most days, the max-out and efficiency effects are not necessarily negligible, because the infrequent times at which these effects are largest are exactly those times when precise estimates are most needed to quantify the effectiveness of Demand Response. The coefficients that describe piecewise-linear relationship between load and temperature, as illustrated in Figure 7, can be estimated by adding yet more terms to the linear regression model. For Office Building A, the resulting temperature-dependence of load is shown in Figure 8. We label the y-axis “excess load” rather than “cooling energy” because there is no certainty that all of the temperature-dependent load represents cooling energy: all we can say is that, statistically speaking, this is the component of load that varies with temperature.

Figure 8: Excess load vs. temperature, for an office building in Martinez, California. Estimates are from a model that includes a piecewise-linear relationship between load and temperature.

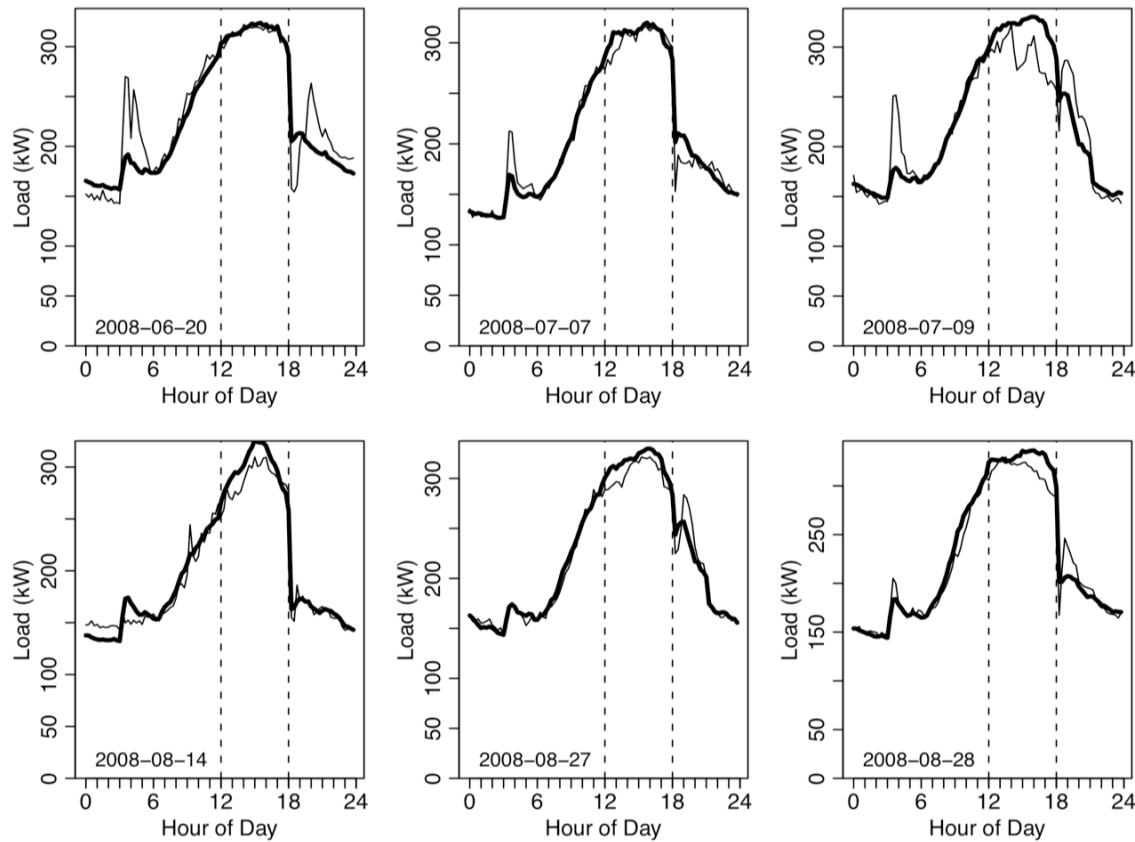


Estimates of Demand Response effectiveness

Linear regression models can also be used to quantify the effectiveness of Demand Response (i.e. reducing load during a specific time period), by incorporating an “indicator variable” that predicts load to differ by a fixed amount in DR periods than in other periods; the regression coefficient associated with the indicator variable provides the best estimate of the load reduction.

Figure 9 shows model predictions and measurements from an office building in Martinez, California, for six DR days in summer, 2008. The thick line shows the prediction *excluding* the DR coefficients, and the thin line shows the measurements. Since the DR coefficients are not included, the thick line shows the predicted load for a *non*-DR day that has the same explanatory variables (temperature, day of week, etc.) as the actual day. An effective demand response would yield an actual load less than the predicted load between 12:00 and 18:00. DR appears to have been ineffective (or perhaps was not even attempted) on June 20 and July 7, but was apparently modestly effective on the other days. Additionally, July 9, August 27, and August 28 all show an apparent rebound following termination of the DR event at 18:00.

Figure 9: Predicted (thick line) and observed (thin line) load in an office building on six DR days in summer 2008. In each case, the prediction ignores the effect of DR. The vertical dotted lines at 12:00 and 18:00 indicate the DR period.



In this building, on an average DR day the load is decreased by 6.4 kW between 12:00 and 15:00, and by 7.5 kW between 15:00 and 18:00, compared to what would be expected for a non-DR day that is otherwise similar. (These represent about a 2% load reduction). The regression also reports a standard error of each of these estimates: for each period, the uncertainty is ± 1.2 kW. However, this uncertainty is an underestimate of the true uncertainty, because it assumes, incorrectly, that the model residuals have no autocorrelation. There are rather involved statistical methods that adjust for autocorrelation of the residuals and thereby give a correct uncertainty, but we do not apply them here; this might be a worthwhile area for future research.

In addition to quantifying the average DR load decrease, it is possible to make a separate estimate for each DR event, by creating a separate indicator variable for each event or by calculating the average residual during each DR period separately. The load reduction averaged over all events, as we have calculated above, can be estimated far more precisely than can the load reduction for any single event, so we recommend estimating the average DR effect rather than the effects for individual events unless the latter are needed for a specific purpose.

Conclusions

Whole-building electric load data can provide useful information about the operation of a building. Data from as short a period as a few weeks, can be used to estimate base load, peak load, duration of high load, and weather sensitivity, and these parameters alone may provide useful information to building operators and researchers.

In addition to results of statistical analyses, simple graphical displays that show load versus time, and facilitate comparison across days, weeks, months, or years, can also be useful. Data collected over a longer duration can be used to monitor changes in the building's operation, and to quantify the effectiveness of operational changes (such as demand-response) or equipment retrofits. In the text, we presented several rules of thumb for creating such displays.

Quantifying some parameters of interest, such as temperature sensitivity, requires fitting a statistical model. In this document, we have used ordinary linear regression, but with special features (time-of-week indicator variables, and piecewise-linear temperature dependence) that are more complicated than casual data analysts are used to. The level of complexity in fitting and interpreting statistical models probably limits the utility of these approaches to researchers, as opposed to building operators, at the moment.

As discussed earlier in this report, additional work is needed to properly account for temporal autocorrelation so as to correctly assign uncertainties to parameter estimates. To convert regression models (or other statistical models) from research tools to tools that will be useful to building operators would require joint efforts of building energy researchers and software developers, to create a robust, easy-to-use interface for performing the analyses and summarizing the results in useful form.

Also, ease of use would be greatly improved by automated selection of parameters that are currently the choice of the data analyst. For instance, for all of the buildings discussed in this report the data analyst first had to look at the load data to determine the starting and ending hours for the period during which the building is in its "occupied" mode, and then enter these hours into the data analysis functions; automating this procedure would improve convenience, but may not be trivial. For instance, hours of operation (or other building parameters) may change with time, perhaps more than once. Such changes can often be readily detected by a person through use of good graphical displays, but automated recognition of these changes may not be so easy.

Finally, statistical models (and the tools to create them) should be developed that can handle the common case that some parameters change while others remain constant. For instance, in the approaches discussed in this report, a change in operation of the building (such as a change in the hours during which the building air is conditioned) requires fitting entirely separate models before and after the change. But some characteristics of the building will not change just because the operating hours shift – for example, the relationship between load and temperature in the middle of the day should be unchanged – so fitting completely independent models is unnecessarily wasteful of data. Developing a single model that can handle the case that some parameters are fixed while others vary with time will require research.

CHAPTER 1:

Graphical approaches to identifying changes and anomalies in electric load

Figure 10: Electric load versus time (at 15-minute intervals) for four buildings in Martinez, California, for June 2008.

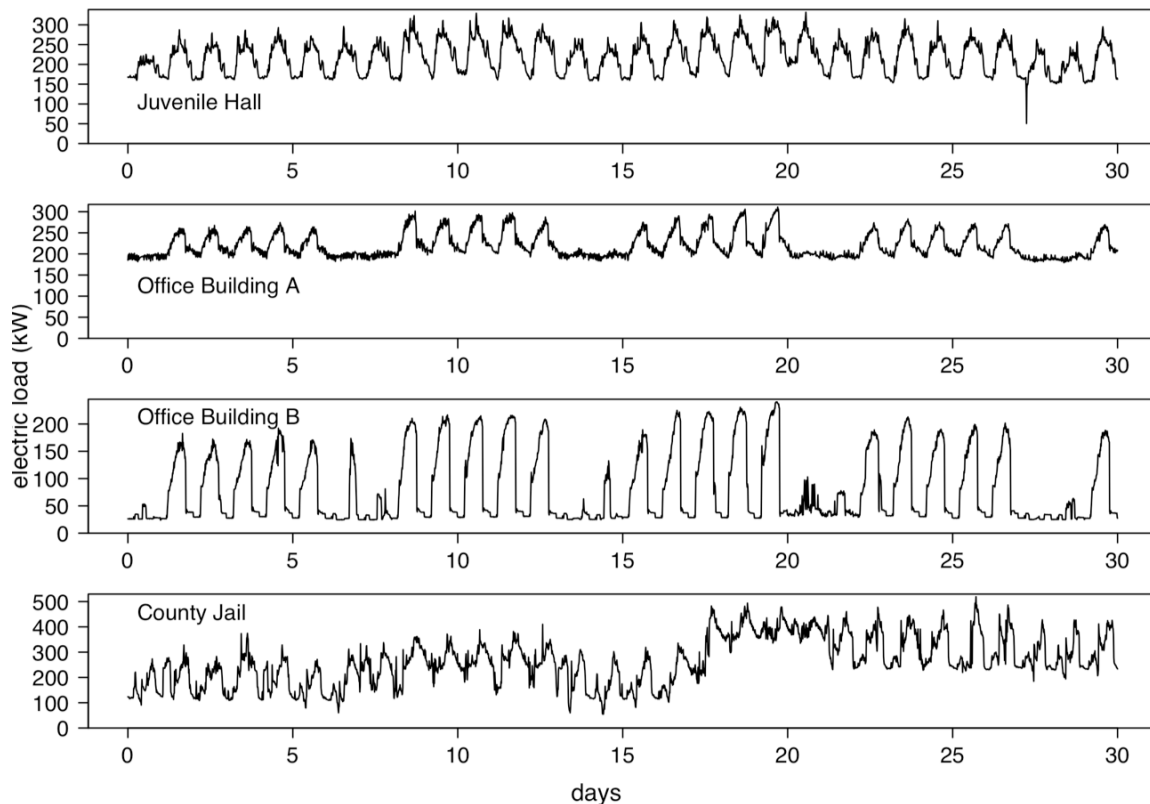


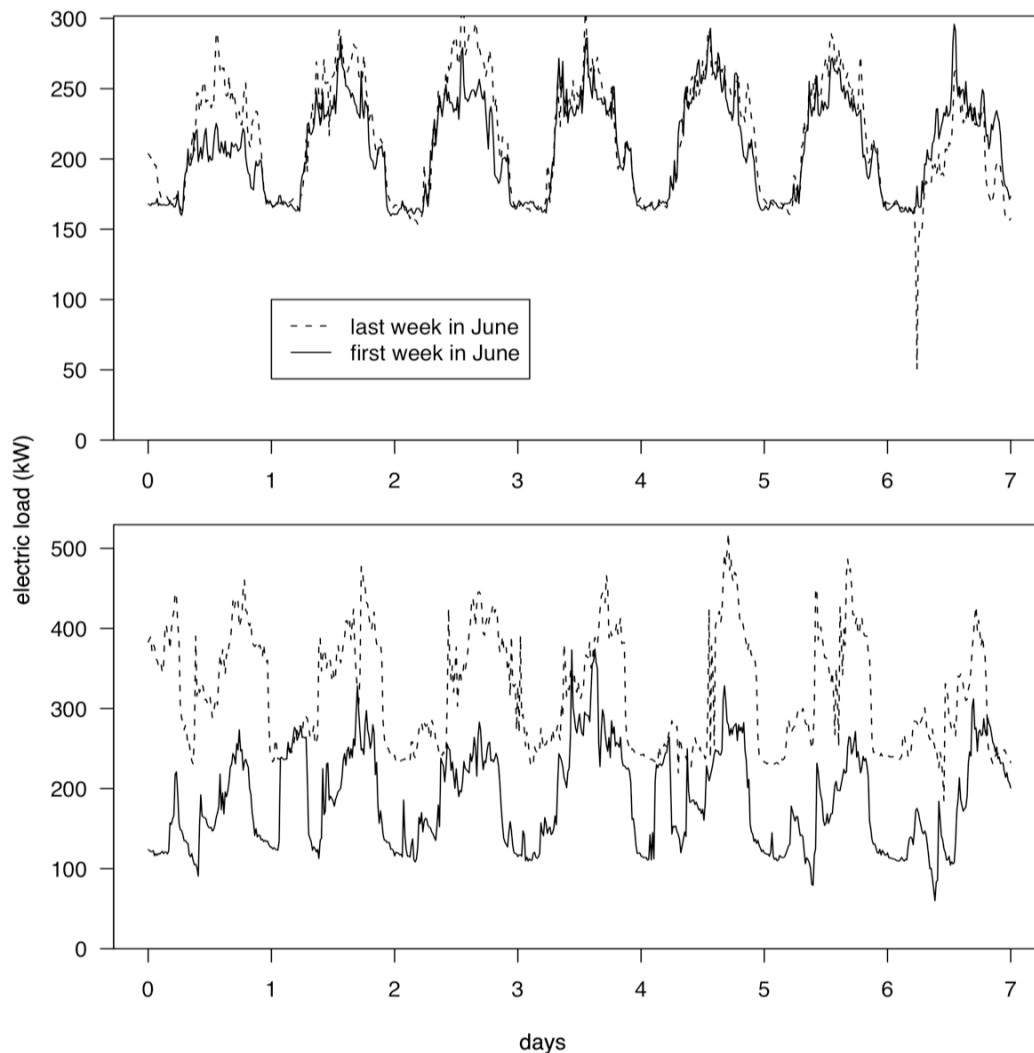
Figure 10 shows electric load versus time in June, 2008, for four buildings in the city of Martinez, California. The first building is the Juvenile Hall (a prison for minors); Office buildings A and B are adjacent office buildings that are occupied by City and County departments; the fourth building is the county jail.

Striking differences between the buildings are immediately apparent. For example:

1. Office Buildings A and B use far less electricity on weekends than on weekdays, whereas the Juvenile Hall shows little weekend / weekday variation.
2. Office Building A has much higher nighttime loads than does Office Building B.
3. The load in the jail varied much more than the others from one week to the next --- for instance, nighttime loads in the second half of the month was almost as high as the daytime load in the first half.
4. Variation from one (fifteen-minute) data point to the next appears to be higher in Juvenile Hall and the jail than in Office Building A or B.

Fully automating the analysis of time-series data such as those in Figure 10 is a daunting task. It is easy to automate summary statistics such as the typical minimum and maximum load, but so many features can vary from one building to the next, one day to the next, and one week or month to the next, that an experienced data analyst looking at the time series will be able to recognize anomalies that automated data summaries are likely to miss.

Figure 11: Electric load for June 1-7 inclusive (solid line) and June 22-29 (dashed line), in Juvenile Hall (top) and county jail (bottom).



Load shapes in most buildings in our dataset show some level of regularity on several widely differing timescales. One weekday often looks much like the next; one week looks much like the next; and in some buildings one year looks much like the next. Simply plotting data from different time periods on the same plot can allow easy comparison of load shapes, and help recognize changes. For instance, Figure 11 shows electric load in Juvenile Hall and in the county jail, for the first and last week in June 2008.

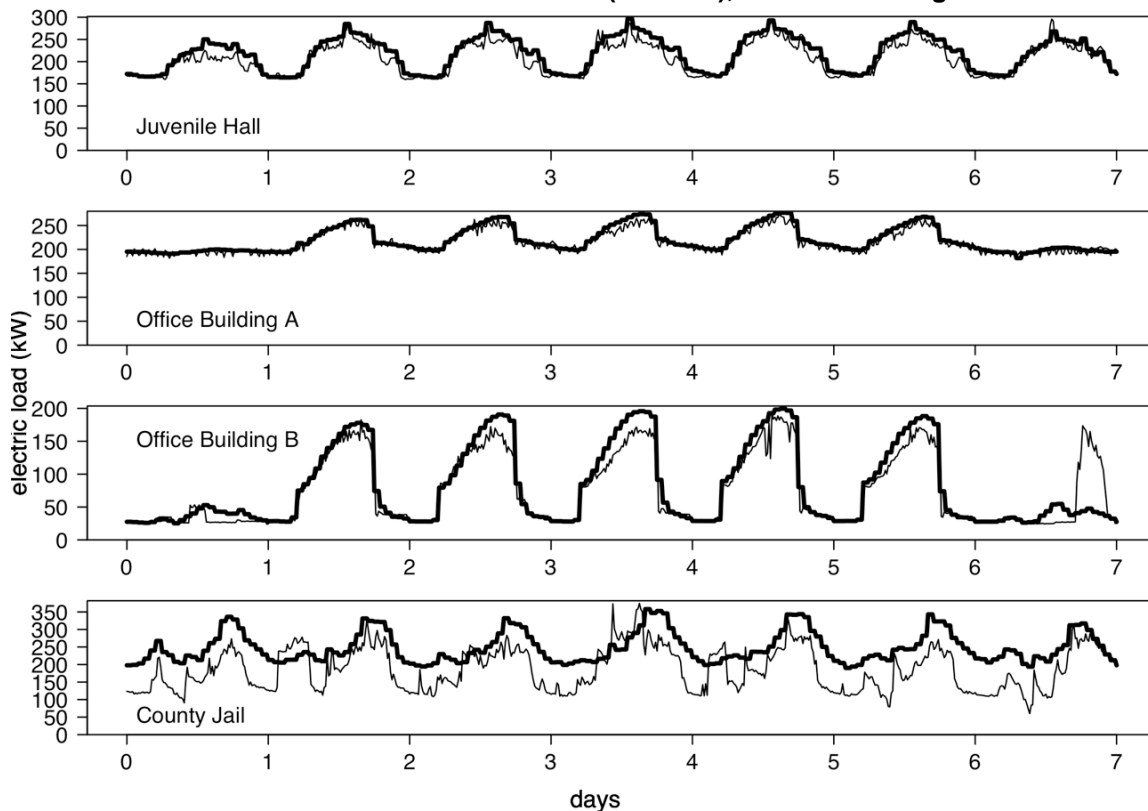
In Juvenile Hall, the load on Sunday (at the left side of the plot) varies substantially from one week to the next: the load throughout the day on Sunday, June 22 is substantially higher than on Sunday, June 1. Monday and Tuesday also show some differences from the first to the last

week – in both cases the load is higher later in the month – but for the rest of the week the two curves coincide well. And the nighttime load is essentially identical for the two weeks.

The situation is quite different for the County Jail (lower plot in Figure 11). As the plot shows, electric load is much higher at the end of the month than at the beginning. Overlaying curves, as in Figure 11, makes it easy to see at least two features that are not apparent from Figure 10. First, the electric load from June 22-29 is shifted upwards by about 150 kW compared with the load from June 1-7, and this shift is approximately uniform throughout the week. Second, in spite of the chaotic appearance of the time series, the temporal behavior of the peaks and valleys is repeatable: with just a few exceptions, most of the significant peaks and valleys match up by time and day of week.

In addition to comparing data from two different time periods, as in Figure 11, it can be helpful to compare data from a single time period to data averaged over many time periods. Figure 12 shows load data from the same four buildings as Figure 10. Plots show the load a single summer week (the first week in June), overlaid with the average data for the same day of week and time of day, where the average is taken over the entire summer. Comparing a single week to a long-term average can help identify anomalous weeks or anomalous data points more effectively than comparing two different time periods, since if two time periods differ it is not clear which is unusual. In contrast, it is clear from **Error! Reference source not found.** that the usage in Office Building B on Saturday evening (the right side of the plot) of the first week in June was atypical. However, if week-to-week variability is very large, as with the County Jail, differences between the average load profile and the profile of a given time period may be less informative, simply because every time period will differ from the average load profile.

Figure 12: Load profile by day of week, averaged over the entire summer of 2008 (thick line) and load from the first week in June (thin line), for four buildings.



There are some helpful general principles when overlaying load shape curves:

1. Displayed in black and white, two or three curves are often the most that can be shown without becoming visually confusing, but with the use of color this can be increased to four or five.
2. The y-axis should always start at zero.
3. Plots that show a time period of a few days, up to about a week, are best for comparing one time period to another. Longer periods require the plot to be too compressed along the time axis.
4. Choose the plot aspect ratio so that major features have a slope of between 30 and 60 degrees up or down, as is the case for Figure 12. Figures 1 and 2 violate this principle slightly: the morning rise in load and the evening fall in load are too steep. As Cleveland (1994) has noted, people are best at comparing slopes near 45 degrees, so very steep or shallow slopes can make it hard to notice asymmetries between the rate of rise and the rate of fall, and to make other useful comparisons. For Figures 1 and 2, the solution would be to show fewer days of data on each plot, but in the present paper we prefer to present more data and are willing to violate the rule of thumb concerning the slopes.

Limitations of Graphical Displays

Although simply overlaying load data from different time periods can be useful for noticing and characterizing changes, there are many limitations to this approach. One problem is that weather-related variation from one period to the next is expected – specifically, energy use is expected to be positively correlated with outdoor temperature (and humidity) during summer,

and to be negatively correlated with outdoor temperature during winter. The question of interest, when comparing data from different weeks or different years, is usually not “is energy consumption higher now than it was during the comparison period,” but rather, “is energy consumption higher now than would be expected, given the weather?”

Another limitation of relying on graphical comparison of load shape in different periods is that some phenomena may be difficult to recognize. For instance, noticing a tendency for load to increase gradually over a long period might be quite difficult, since this small trend will often be superimposed on seasonal variation and other features. In this example, overlaying data separated by a year will reveal that load is higher now than a year ago, but will not reveal whether the change was gradual or abrupt. As we discuss below, this is a problem for automated analysis methods and not just for graphical approaches.

Finally, and most importantly, graphical approaches to comparing load shapes are useful only when people are able to devote time to using them: it takes time and effort to generate a plot, to look at it carefully, and to interpret it. Automated analysis methods can potentially reduce the effort required to identify problems or opportunities for energy savings.

CHAPTER 2:

Parameters that Describe Load Shape

Typical office building load shape in summer

As the examples above illustrate, building load shapes can vary widely. However, many office building load shapes share some characteristics. Figure 13 shows a stylized summer office load shape. Its features are explained as follows.

In summer, outdoor temperatures are high even at night. Since the setpoint temperature is usually higher at night than during the day, the building is warm in the morning when the setpoint temperature is returned to its daytime value. The HVAC system immediately turns on at high power, to cool the building to the daytime setpoint temperature, resulting in a sudden spike in load, a phenomenon known as the “morning catchup” or “morning rebound.” The duration of this catchup spike depends on the temperature inside the building before the setpoint changes. There may be no spike at all, if the indoor temperature is already below the daytime setpoint.

While the building systems are returning the building to its “occupied” mode, building occupants are arriving, so load increases as they begin to use lights, computers, and other office equipment. Outdoor temperature is also increasing, which increases the cooling load. The increasing occupancy and cooling load cause the load to increase steadily for typically two to three hours, which we call the “morning ramp-up.”

Eventually the load begins to level off, though it is still affected by outdoor temperature (and insolation) so the load often peaks in the early afternoon. In the late afternoon or early evening, people begin to leave, the outdoor air cools, the sun sets, and the temperature setpoint is returned to its (higher) nighttime value, so load drops precipitously. In most office buildings, the evening setback effect is much sharper – that is, takes place much more quickly – than the morning ramp-up. However, in some buildings some excess over the base load persists into the evening, due to cleaning crew and to workers staying late, leading to a ramp-down or “shoulder” before the load again falls to its late-night level, or “base load.”

Figure 13: Stylized load shape for a typical office building in summer.

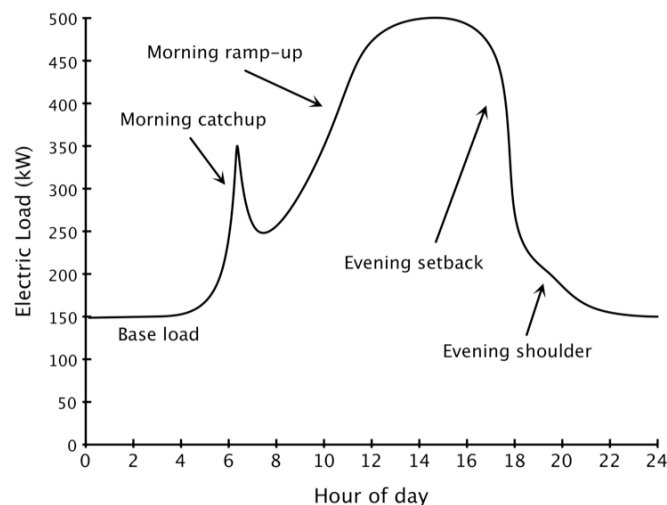
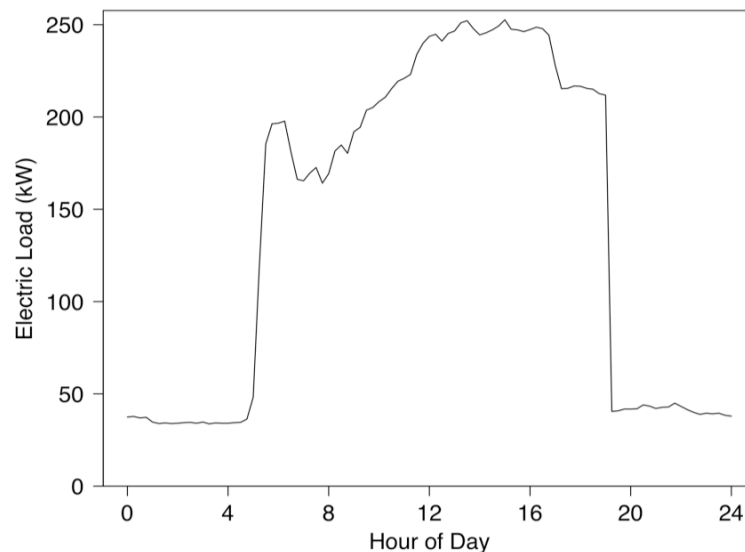


Figure 14 shows the electric load for a government office building in Martinez, California, for one day in July 2008. The base load, morning catchup, morning ramp-up, and evening setback are clearly visible, but there is no evident evening shoulder. The evening setback in this building takes place rather late, around 7:30 PM (hour 19.5), which is three hours after a sharp drop in load around 4:30 PM (hour 16.5). It is possible that most people leave the building at 4:30 PM, but that most of the building's lights and cooling systems remain on until 7:30 PM.

Figure 14: Actual load versus time for an office building in Martinez, California in July 2008.



Summary Statistics Without Weather Adjustment

As with any time series, load curves invite the calculation of a wide variety of summary statistics. It is useful to distinguish between two kinds of summary statistics: those that summarize various aspects of the load and its variability, and those that summarize the amount of load or load variability that is not related to temperature or humidity. For instance, consider a building's base load. The base load might vary from month to month, but this variation might be almost perfectly predicted by the nighttime outdoor temperature. For some purposes, a data analyst might care only about the variability, whatever the cause; for others, they might be interested only in the amount of variability that is not related to weather.

It is common in load analysis to examine values such as mean, median, maximum, and minimum values; for instance, what is the mean load at 2:00 PM on a Tuesday, and what are the highest and lowest loads observed at that time of the week. Depending on the application, maximum and minimum values may not be the best choice because by their very nature they summarize only the extreme events, which may differ qualitatively from events that are actually of interest. This is particularly true of the minimum load: an extreme event, such as shut-off of all or part of a building's electricity for service or emergency, can lead to a load that is so anomalous that it has little or no relevance for quantifying the effect of normal building operations. Having one or two events like this will have little effect on the mean load, and perhaps no effect on the median, but these events will define the minimum (and other anomalous events may define the maximum). We recommend using a low and high quantile, such as the 2nd and 98th percentiles, rather than the minimum or maximum.

Figure 15: Sketch indicating five parameters that we recommend as a minimum set to characterize load shape. The value of each parameter can be calculated for each day; these values can then be summarized (e.g. average peak load, standard deviation of peak load).

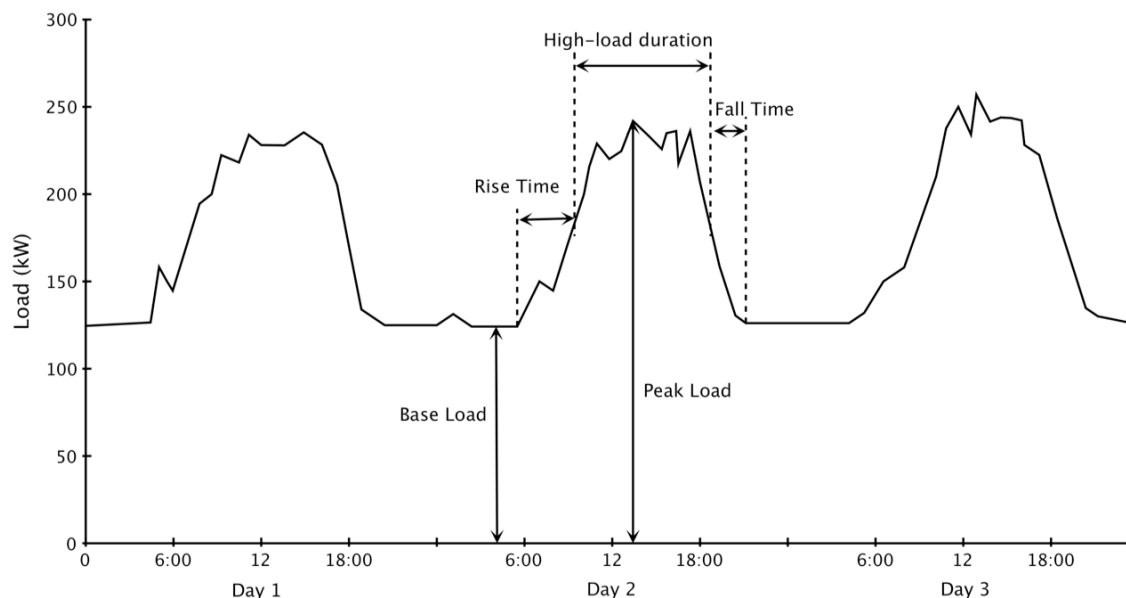


Figure 15 shows several parameters that can be used to summarize load shape. Each day, the base load and peak load are calculated, as well as three time intervals. The first time interval quantifies how long it takes for the load to increase halfway to the peak load (i.e. to reach a load closer to the maximum load than to the base load); we call this the “rise time.” The second quantifies how long the load stays above the halfway mark; we call this the “high-load duration.” And the third, the “fall time,” quantifies how long it takes the load to fall from the halfway point back to the base load. Additionally, we summarize the variability of the load during the high-load period, by calculating the standard deviation of the load during that time.

Some buildings, or at least building data, have a peak load – the very highest 15-minute data point – that is substantially higher than any other data point. In some cases we are interested in that highest data point, but more often we are interested in the behavior of the building at slightly longer timescales. For the latter purpose, we define a daily “near-peak load” to be the 97.5th percentile load, i.e. a value that is greater than 97.5% of the load measurements that day. In many buildings and for many purposes, this value is both more stable (i.e. exhibits less day-to-day variability) and more relevant than the peak value. Similarly, we define a daily “near-base load” to be the 2.5th percentile load.

The schematic load shape shown in Figure 13 or Figure 15 leaves little doubt as to when the load begins to increase from the base load, or returns to the base load; this is the case for many real buildings as well, as Figure 14 illustrates. However, in some buildings the situation is more complicated (see data from the county jail in Figure 11, for example). Indeed, some buildings, such as industrial facilities, may have load shapes that look completely unlike the schematic load shape in Figure 15 and may require different descriptive parameters.

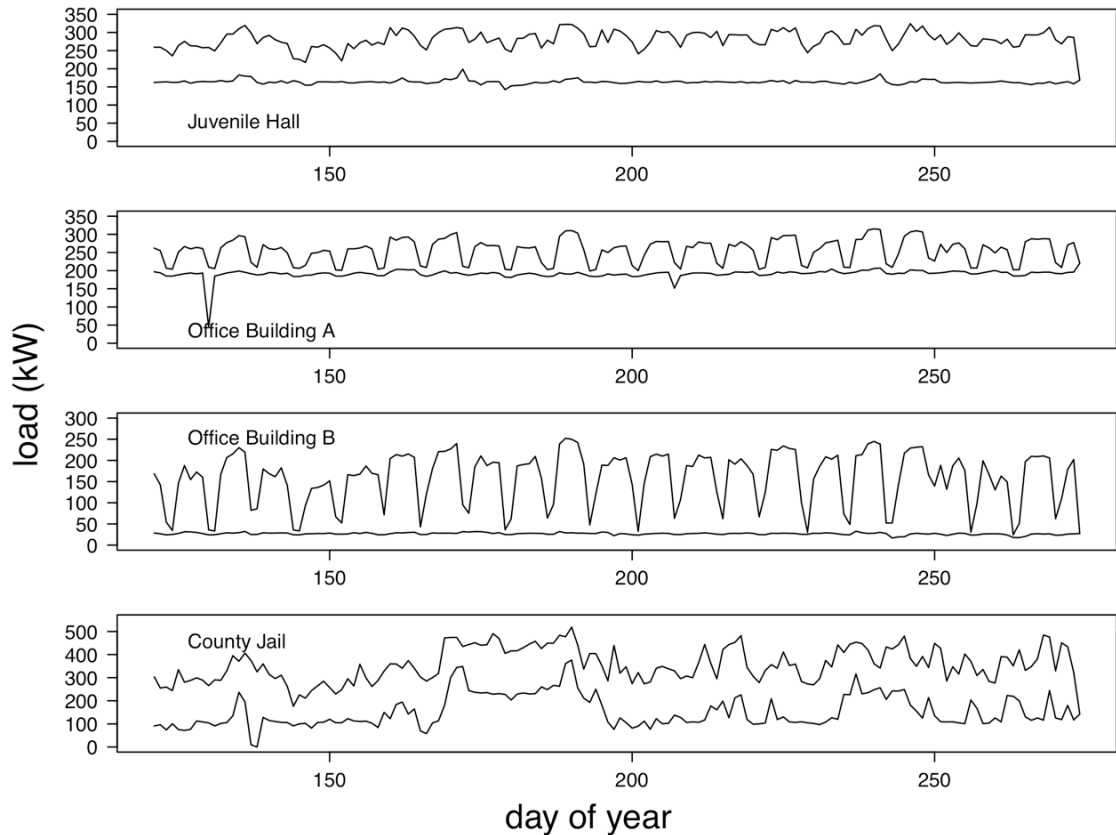
In practice, it is hard to find a definition of “rise time” and “fall time” that yield consistent, easily interpretable results. We have experimented with various definitions. The most satisfactory that we have found thus far is to define the rise time as the period between the latest time in the morning that the load is close to the base load, and the earliest time during

each day when the load is more than halfway to the 97.5 percentile load for the day. By “close to the base load”, we mean that the load is less than: $\text{base load} + 0.05 * (\text{near-peak load} - \text{near-base load})$. This definition has reasonable behavior if the load shape is something like Figure 15, but it does not produce useful numbers in some cases, such as load profiles that do not vary substantially over time, or very “spiky” load profiles that have multiple extreme maxima and minima during each day. Determining a generalized, automated way to usefully classify load shapes by rise time (or fall time) might be a fruitful area of future research.

Table 2: Parameters that characterize electric load shape (for a single day) in a typical office or retail building.

Parameter	Definition
Near-Peak load	97.5 percentile of daily load
Near-Base load	2.5 percentile of daily load
High-load duration	Duration for which load is closer to near-peak load than to base.
Rise time	Duration for load to go from base load to the start of the high-load period.
Fall time	Duration for load to go from the end of the high-load period to the base load.

Figure 16: Nighttime near-base load (lower curve) and daytime near-peak load (upper curve) by day in the same four buildings whose load is shown in Figure 10 for the spring, summer, and fall of 2008. In contrast to Figure 10, this figure shows only the daily values, not the variation within the day.



The statistics listed above can be calculated and summarized for various timescales. Table 3 shows several statistics summarized over the spring, summer, and fall of 2008, for the four buildings discussed above. Statistics are calculated separately for weekdays and weekends. As the table shows, for these buildings, there is little variation in base load between weekends and weekdays. But the Near-Peak Load column shows that even in Juvenile Hall and the Jail, the near-peak load is lower on weekends than on weekdays, a fact that is not immediately evident from Figure 16. Also, the coefficient of variation (COV, defined as the standard deviation divided by the mean) of the daily near-peak loads is much higher in Office Building B and the jail than in the other buildings. The COV is a measure of variability relative to the mean: a COV of 0.1 means typical variation is about $\pm 10\%$ of the mean value.

Finally, whereas the other columns are calculated from the daily near-base and near-peak loads alone, and thereby summarize the day-to-day variation in near-peak and near-base load, the last column summarizes the amount of load variation *within* a given day: the COV of the load during the high-load period is calculated for each day, and the median of these values is recorded in the table. If the load typically jumped to the near-peak value, stayed there for most of the day, and then suddenly dropped back down – i.e. if the load profile had a very flat top – then the value in this column would be 0. If, instead, the load varies substantially with time, even within the high-load period, then this value will be larger. A value of 0.25, for example, would mean that the variation with time, during a typical high-load period, is 25% of the mean load during the period. The large weekday value of 0.66 for Office Building B on weekdays

indicates that even during the high-load period, this building's load changes a great deal throughout the day.

Table 3: Median daily near-base load, median daily near-peak load, Coefficient Of Variation (COV) of daily peak load, and median COV of load during the high-load period (i.e. the period when load is closer to the peak than to the base), for four buildings in Martinez, California, for days in spring-fall 2008. Tabulated values are for weekday, weekend. Near-base COV is not shown to save space.

	Daily			Within-day
	Near-Base Load (kW) Median	Near-Peak Load (kW) Median	Near-Peak Load COV	High-load COV Median
Juvenile Hall	163, 163	294, 264	0.08, 0.07	0.18, 0.16
Office Bldg A	194, 186	271, 207	0.07, 0.04	0.11, 0.03
Office Bldg B	28, 25	198, 64	0.18, 0.51	0.66, 0.36
County Jail	121, 132	350, 329	0.20, 0.22	0.27, 0.22

Adjusting for Weather and Other Factors

There are at least two types of load prediction: (1) load *forecasts* are predictions of future energy use, and may be used for advanced building control, DR response planning, and other purposes; and (2) *post-facto predictions* can be made for comparison with historical data to identify anomalies and faults, to determine the effect of changes in building operation, to verify retrofit effectiveness, to determine the temperature sensitivity of a building's energy consumption, or for other purposes. Load forecasting is more challenging than post facto prediction, because forecasting is necessarily based entirely on data that are available prior to the time period being forecast. For instance, consider predicting peak energy consumption for the first day of the spring in which cooling is needed: the forecast must be made without any recent data that include cooling load, whereas a post-facto analysis can include data collected later in the spring. In this report, we focus primarily on post-facto analysis rather than on forecasts.

There are two types of approach to predicting electric load based on explanatory variables: (1) find a mathematical relationship between explanatory variables and load, and apply this relationship to make predictions; or (2) predict the load to be the same as it was when explanatory variables had values close to those at the time period of interest. Each of these types of approach leads to several variations, discussed below.

Several methods have been applied by researchers or as part of Energy Information Systems (EIS) to determine mathematical relationships that predict load. Granderson et al. (2009) noted that the following methods are in current use:

1. Linear regression of load on outdoor temperature and humidity, with indicator variables for time of day or time of week, yields a model that assumes that there is a linear relationship between each explanatory variable and the load. This assumption is justified in many cases. For instance, the rate at which heat is conducted into a building is proportional to the indoor-outdoor temperature difference. Since the cooling system must eject heat at the rate that it enters the building, in order to maintain a constant

indoor temperature, the cooling energy needed is nearly linear in outdoor temperature over a wide temperature range. However, there are cases for which the linearity assumption are not justified. For example, once the cooling system is operating at its maximum the cooling load cannot increase, no matter how high the outdoor temperature. Other phenomena, such as variation of chiller efficiency with load, also lead to a more complicated relationship between temperature and electric load. As we illustrate below, even rather complicated relationships can be handled with linear models.

2. Non-linear models are similar to linear regression models, but can allow a more complicated relationship between input variables and output variables. Such models may be more accurate than linear models at extremes of temperature (when heating or cooling systems are working at maximum load) and when outdoor temperatures are mild (when heating and cooling systems may be turned off, so that electric load is independent of outdoor temperature over some temperature range). A type of nonlinear model that is attractive for load shape analysis is “locally weighted scatterplot smoother” (lowess) and a closely related approach called loess, which we discuss below in more detail.
3. A neural network model can create a very complicated mathematical relationship between explanatory variables and the predicted load. Input variables such as outdoor temperature, humidity, etc., are mathematically processed to create a potentially large number of secondary, or “hidden”, values. These hidden values are then processed to generate a (usually small) number of output values, such as predicted load. The mathematical functions that process the input values and the hidden values have adjustable parameters known as weights, so that the effect of every input value on every hidden value is adjustable, as is the effect of every hidden value on every output value. Neural networks “learn” by adjusting the weights so that the outputs are as close as possible to their desired values, for a large set of “training” data. Neural networks can generate a rather complicated mathematical relationship between explanatory variables and predicted load, but can behave poorly when conditions differ substantially from all of the training data. Neural nets are also very hard to analyze: they are essentially “black boxes” from the user point of view, in the sense that it can be very difficult to understand or explain the relationship between input variables and predictions.
4. For each data point – each five-minute, fifteen-minute, or one-hour interval, for example – break the parameter space of explanatory variables into bins. For example, temperature might be binned into five-degree intervals, time of week into hour of the week, and relative humidity into five-percent intervals. Given the values of the explanatory variables, find all of the historical data that fall in the same bin. Predict the load to be the average (or median) of the data in the bin. Advantages of the bin approach are that it is simple to explain and understand, and works well when only a few input variables are important. Bin models can only handle a small number of input variables (because if there are too many variables, explanatory variables may not fall in the same bin as any previous data), and they may be unable to predict, or may provide inaccurate predictions for conditions that have occurred rarely or never before.
5. Another approach applies the same basic principle as the bin method –the predicted energy consumption is the average consumption during similar periods – but avoids many of the problems of the bin method. It creates a metric to describe the degree of similarity between the explanatory variables at the time that is being forecast, and conditions at other times. There is no limit on the number of input variables that can be used effectively, but, as with bin-based methods, results can suffer if the current conditions have rarely or never been encountered before, so that even the nearest neighbors are distant.

Variants of the above approaches are also possible. For instance, instead of a single linear model that applies through the entire year, two models could be constructed, one for cooling season and one for heating season. Or, instead of predicting the load to be the average load at times when conditions were similar to the given conditions, a subset of moderately similar conditions can be identified and used to create a linear model that applies only to a portion of parameter space. There are other possibilities as well.

Any of the load prediction methods mentioned above can be used for several purposes, including fault detection (by finding times when the building isn't behaving as it has behaved in the past) and quantifying the effectiveness of Demand Response (DR). In each case, the predicted load is compared to the actual load.

In this report, we discuss various linear regression models because – when constructed appropriately – they provide a good fit to load data in most buildings; their results are easy to interpret; they are easy to modify; and they present modest computational burden. This latter point is important, even with today's computers, when working with load data collected at 15-minute intervals: one year of data comprises about 35,000 data points for each variable in the model, which makes the full analysis large enough that some software products can struggle with it or fail altogether.

Although models can be created to use many explanatory variables, such as wind speed and direction, building occupancy data, and so on, we focus on predicted loads that depend only on historical load data, outdoor air temperature, and humidity. These are the only data that are available in many cases, and equally importantly, these few variables are sufficient to explain most of the load variability in most buildings. By restricting ourselves to these variables, we create models that can be applied to almost any building load dataset.

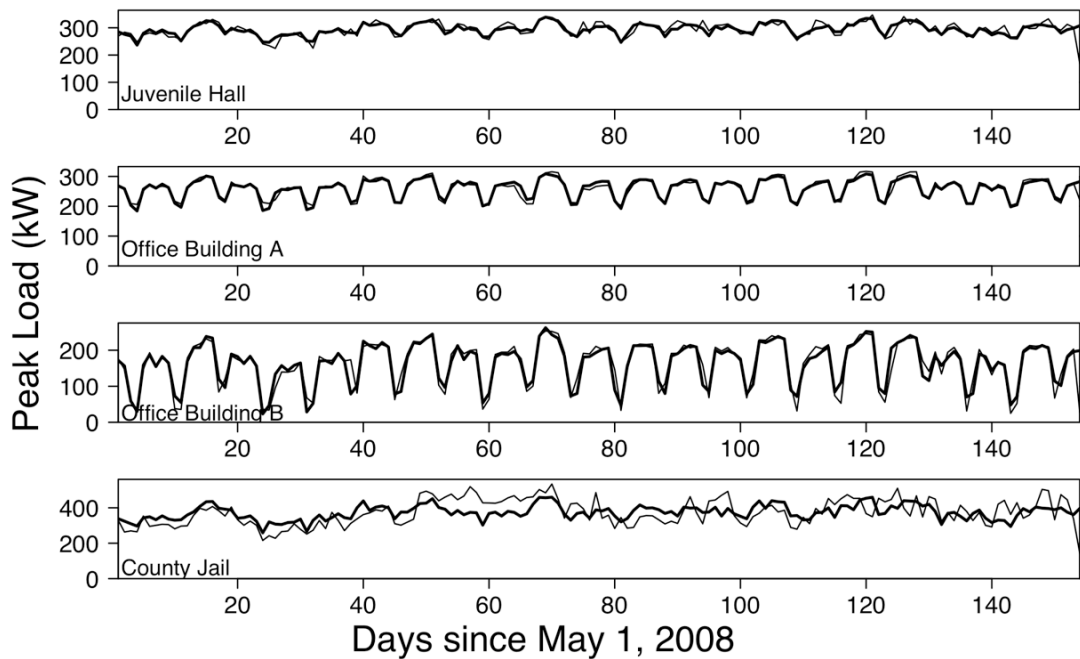
We begin by considering a specialized application of load shape analysis: predicting only the daily peak load. Even if peak load were the only quantity of interest the best approach is to model the load throughout the day; we discuss the prediction of load as a function of time of day in a later section. However, peak load is important from the standpoint of DR. Moreover, restricting our discussion to peak load allows us to introduce the terminology and concepts of regression models in a simple setting, while discussing a very simple model that is adequate for many buildings. Consider the model described by **Equation 2**, where \hat{y} is the predicted peak load for a given day of the week and a given temperature:

$$\hat{y}(\text{day of week}, T(t)) = \beta_{\text{day of week}} + \beta_T T(t) \quad \text{Equation 2}$$

where the β values are regression coefficients. We make the time-dependence of temperature explicit for convenience in future equations. In this model, each day of the week has its own typical peak load, and the load increases proportionally with temperature.

For three buildings featured in Figure 10 this model works very well for predicting summer daily peak load, but for the jail the fit is rather poor. Figure 17 shows peak load for each day in summer, as predicted (thick line) and measured (thin line) for the four buildings.

Figure 17: Predictions (thick line) and observations (thin line) of daily peak electric load for several months beginning May 1, 2008.



As briefly discussed in the list of analysis methods a few pages ago, there are advantages and disadvantages to using linear regression to analyze load data. In most cases we have investigated, linear models perform well when the goal is to predict hourly load or daily peak load from time of day, day of week, temperature, and humidity. It is possible that linear models would not be satisfactory for estimating other, possibly nonlinear effects such as the effect of wind direction, wind speed, sun angle, building occupancy, and so on, but including such predictive variables is unusual in practice, in part because the necessary location-specific data are rarely available. Linear regression will be our main analysis tool in the present report.

Summarizing errors in predicted electric load

A standard way of summarizing the degree of fit between predictions and data is to quote the value of r (the correlation) or R^2 (the squared correlation, which is the ratio of variance in the data that is explained by the model to total variance in the data). For either of these, a value near 1 means good agreement between prediction and data, and a value near 0 means poor agreement between prediction and data. However, these are probably not particularly relevant summary statistics for evaluating the usefulness of models for DR. The problem is that the r - and R^2 values depend on the variance of the data, including (in this case) the weekday and weekend variability. This is illustrated by comparing Juvenile Hall and Building B: Building B is much more variable, because of its large “weekend effect” (load is much lower on weekends than on weekdays). This large amount of variability is captured by the model, since it includes a different base load for each day of the week, so the model explains most of the variance in Building B’s load. In contrast, Juvenile Hall is much less variable from day to day, and random (unexplained) variability is a larger fraction of the total variability than is the case in Building B. As a result, R^2 is higher in Building B than in Juvenile Hall, as shown in Table 4. The other statistics in the paper, Standard Error and GMRAE, will be discussed below.

Table 4: Measures of model fit for linear regression prediction of daily peak load in Building A and Building C, in summer 2008. The model includes only day of week and outdoor temperature.

Building	Standard error (kW)	R ²	GMRAE (%)
Juv. Hall	15.3	0.68	2
A	9.9	0.92	1
B	21.9	0.88	5
Jail	62.6	0.31	11

For DR purposes, summaries of model fit other than r or R^2 may be more useful, either for comparing how different models behave or for comparing how a single model behaves across different buildings. The standard error is one simple metric that can be useful: this is the standard deviation of the residuals of the regression model, i.e. the standard deviation of the error in the predictions. In Table 4 we see that the linear model typically predicts the peak load in Juvenile Hall within about 15 kW, and predicts the peak load in Building B within about 22 kW. In these buildings, which are at least roughly comparable in terms of typical power consumption, simply comparing the standard errors makes sense. But to compare model fit for buildings that differ widely in load, it is useful to have a metric that takes the typical load into account.

If $y(t)$ is the load at time t , and $\hat{y}(t)$ is the predicted load at time t , then the relative absolute error (RAE) at time t is the magnitude of the error in the prediction as a fraction of the actual value:

$$RAE(t) = \frac{|y(t) - \hat{y}(t)|}{y(t)} \quad \text{Equation 3}$$

RAE can be usefully expressed as a percentage by multiplying by 100%: an RAE of 0.11, or 11%, means that the error was 11% of the load.

The geometric mean (GM) of the RAE, often abbreviated GMRAE, summarizes the RAE over all of the measurement times. If the RAEs are not highly variable, this single number can be thought of as the typical error. The GMRAE of n data points is given by

$$GMRAE = \sqrt[n]{RAE_1 RAE_2 \cdots RAE_n} \quad \text{Equation 4}$$

In practice, one calculates the RAE by taking the mean of the logarithms of the RAEs, and then exponentiating. This is mathematically equivalent to Equation 4.

As with RAE, GMRAE can be expressed as a percentage by multiplying by 100%. GMRAE, expressed as a percentage, is shown in Table 4. The simple regression model typically predicts the daily peak load in Juvenile Hall to within about 2% of the actual value; in Building B the error is about 5%. Some researchers prefer to use the Mean Absolute Percentage Error (MAPE), which is the arithmetic (rather than geometric) mean of the RAE values, expressed as a percentage.

CHAPTER 3:

Linear regression models for load shape

As discussed previously, one of the common approaches for predicting electric load (either at a given moment, or averaged over an interval) is to use linear regression modeling. Linear models have several desirable characteristics: they are easily implemented in standard analysis programs, and results are easily summarized and easily understood.

A linear model for load at a given time can be constructed along the same lines as **Equation 2**. Divide a week into intervals (indexed by i); for instance, if electric load data are available for 15-minute intervals, the first interval is from midnight to 12:15 on Sunday morning, the second interval is from 12:15 to 12:30, and so on. A different regression coefficient for each period allows each period to have a different predicted load. Additionally, we expect that for all intervals there is a temperature effect: when the outdoor temperature is high, cooling load will increase with temperature, and when the outdoor temperature is low, heating load will increase as temperature decreases. For some range of moderate temperatures, the load may be insensitive to temperature because neither cooling nor heating is needed. The outdoor temperature at which cooling (or heating) is needed can be estimated if there is enough data from a building, although this estimation adds complexity. For example, load may increase as outdoor temperature increases when the temperature is over 65 F, and may increase as outdoor temperature decreases when the temperature is under 50 F.

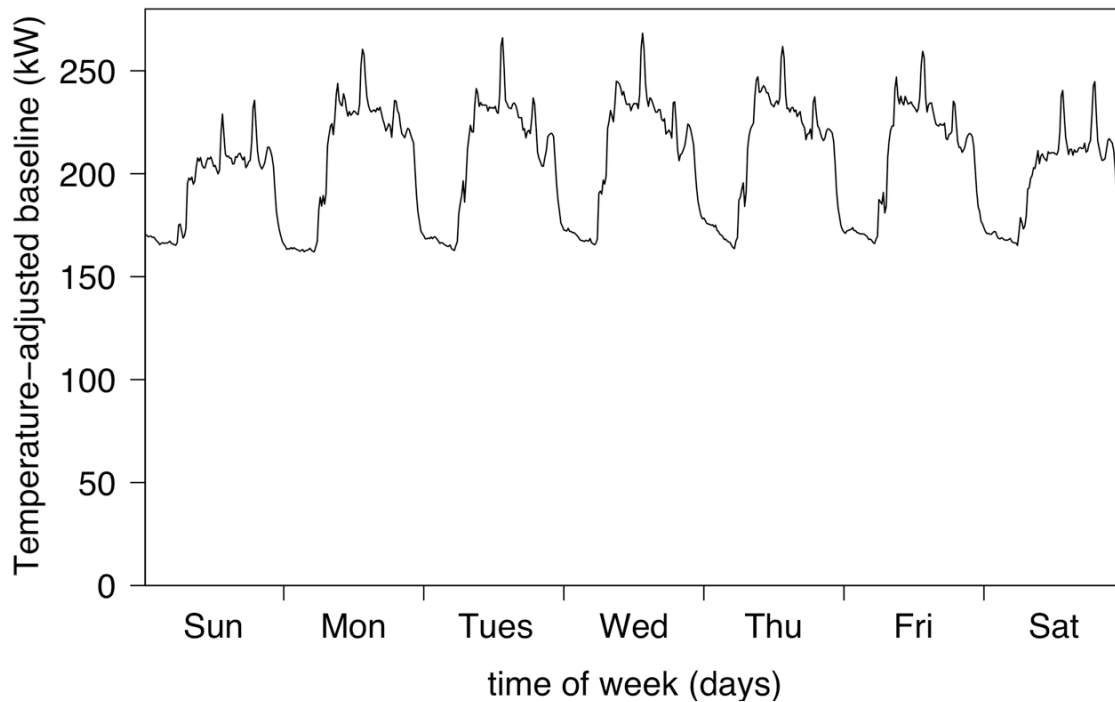
The resulting model is implemented by fitting **Equation 5**, where i indexes the time interval. β_i is the regression coefficient of an indicator variable that indicates what time interval during the week contains the i th time interval. For example, the $i=1$ indicator variable takes the value 1 if the time interval is for a Sunday from midnight to 12:15 a.m., and takes a value 0 otherwise. The $i=2$ interval takes the value 1 if the time interval is for a Sunday from 12:15 a.m. to 12:30 a.m., and zero otherwise. And so on.

$$\hat{y}(t, T(t)) = \beta_i + \beta_{T_{\text{over}}} T_{\text{over}}(t) + \beta_{T_{\text{under}}} T_{\text{under}}(t) \quad \text{Equation 5}$$

T_{over} is defined to have a value of 0 when the outdoor temperature T is under 65 F, and is equal to $T - 65$ F when $T > 65$ F. T_{under} has a value of 0 for $T > 50$ F, and a value of $50 - T$ when $T < 50$ F. (For the time period we consider below, the outdoor temperature was rarely below 50 F, so we will not discuss the low-temperature behavior).

Fitting **Equation 5** results in an estimate of 674 regression coefficients: one coefficient for each 15-minute interval during the week (the β_i values), plus the two temperature coefficients. The coefficient $\beta_{T_{\text{over}}}$ quantifies the additional electricity consumption for each degree that the outdoor temperature exceeds 65 F, and the coefficient $\beta_{T_{\text{under}}}$ quantifies the additional consumption for each degree that the outdoor temperature is below 50 F. In this model, the β_i values have a simple interpretation: they predict the electric usage during the specified time interval, if the outdoor temperature is between 50 F and 65 F. We refer to β_i as the “temperature-adjusted baseline.”

Figure 18: Temperature-adjusted baseline for Juvenile Hall, for each 15-minute interval throughout the week.



We fit the data from Juvenile Hall, from May through September 2008, using the model defined by **Equation 5**. Figure 18 shows the temperature-adjusted baseline. This is simply a plot of the β_i regression coefficients from **Equation 5**, in order by time interval i . We might expect that, at least during the summer, there are times when the air conditioning compressors are running and times when they are not, so that on any given day there will be sharp peaks in load. But these peaks would not necessarily be expected to occur at the same time each day. So it is somewhat surprising that, in Juvenile Hall, the temperature-adjusted base load has one or two well-defined, high, sharp (i.e. short-duration) peaks each day, including a prominent one at 1:30 PM. The load must have a peak almost every day at these precise times, or they would not show up in the temperature-adjusted baseline data, which averages over many weeks of data.

For Juvenile Hall, the estimated value of $\beta_{T_{over}}$ is 2.0 kW/F. If the outdoor temperature is 85 F, then $T_{over} = 20$ F and the model predicts a 40 kW increase in load. During the summer, daytime temperatures at the building are often quite high, while nighttime temperatures are mild, so the result of the temperature sensitivity of the load is to increase the daytime load while leaving nighttime load nearly unchanged. The result can be seen in Figure 19, which compares the temperature-adjusted base load to the model predictions that include temperature, for a few days during a warm week in June. The daytime temperature reached as high as 97 F during this period, but nighttime temperatures were in the 50s and 60s. In essence, the “temperature-adjusted baseline” determines the average electric load profile if the building always experienced a mild outdoor temperature, whereas the model prediction incorporates the effect of temperature.

Figure 19: Temperature-adjusted baseline, and the load predicted by the model described in Equation 5, for Juvenile Hall, Sunday-Wednesday of the second week of June 2008. The thin line is the same as that for Sunday-Wednesday in Figure 18. The predicted effect of high daytime temperatures is easily seen in the fact that the model prediction (which includes temperature) is well above the baseline during the daytime, as expected.

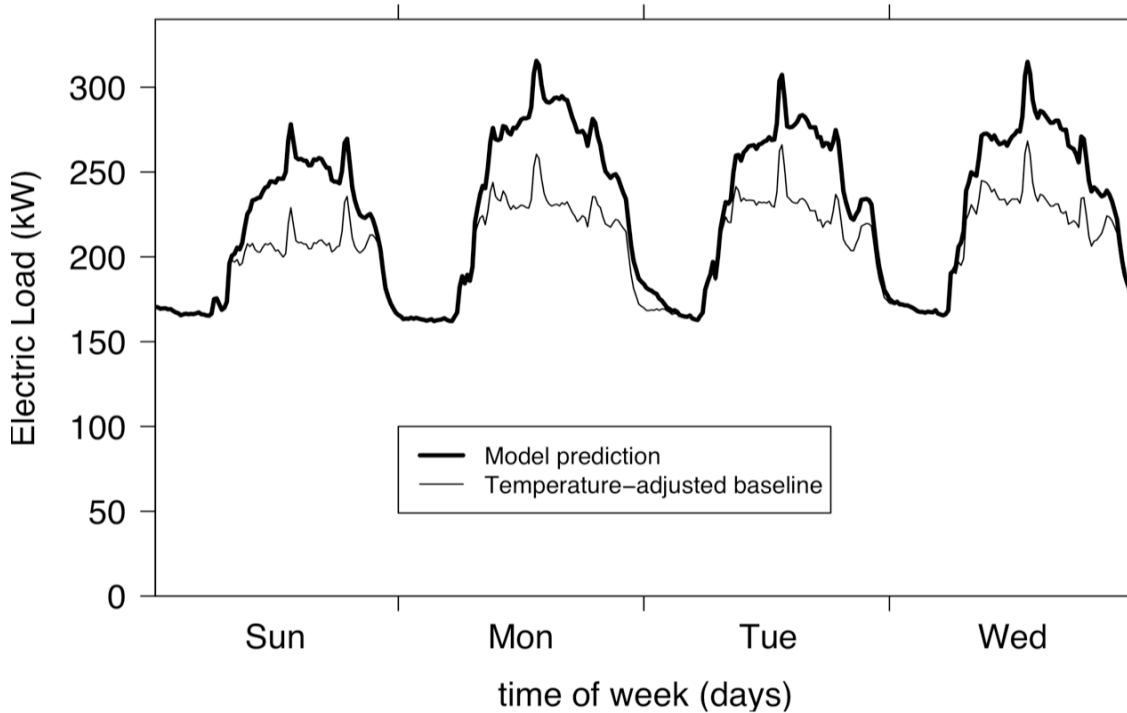


Figure 20 shows the model predictions (same as the thick line in Figure 19) along with the actual load data from a portion of the second week of June 2008. The difference between measurements and predictions, as a function of time, is shown at the bottom of the figure. The data (thin line) show a prominent peak at 1:30 PM each day, which indeed occurs almost every day at that time. But the data also show sharp peaks and valleys throughout the day that are not predicted from the simple model. Perhaps some of these reflect the behavior of the cooling system: it is predictable that on a hot day the cooling energy will be large, but the exact temporal behavior may be unpredictable at the 15-minute level.

Figure 20: Predicted (thick line) and measured (thin line) electric load in Juvenile Hall for four days in June 2008. Data and predictions are for 15-minute intervals. Residuals, i.e. the difference between observed and predicted load, are also shown.

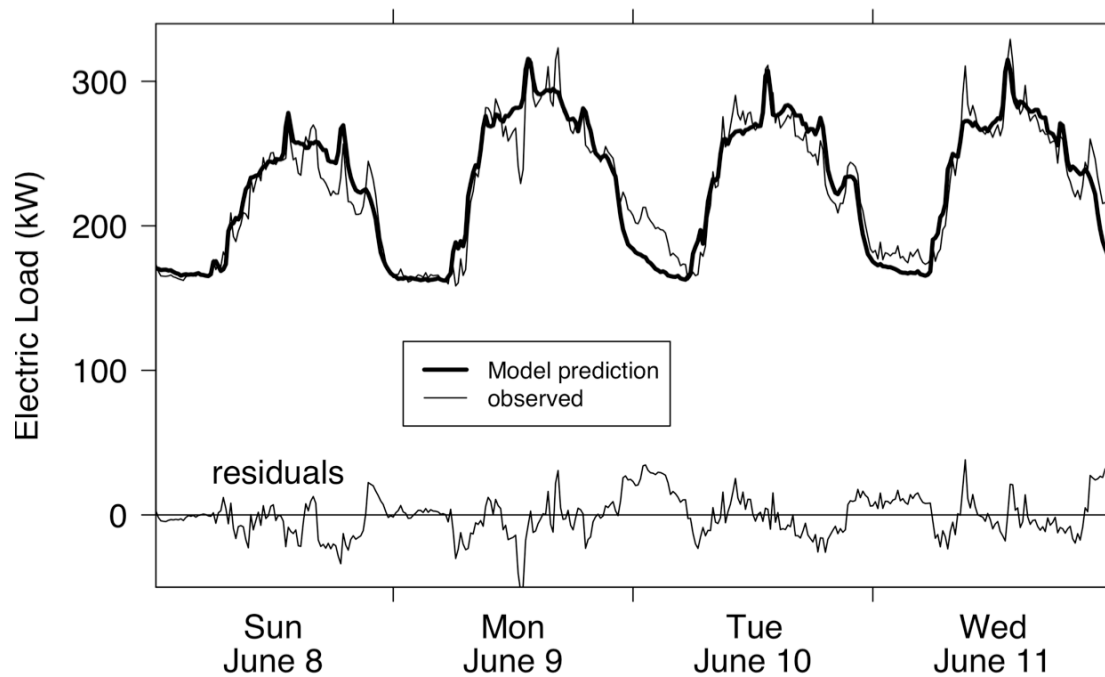
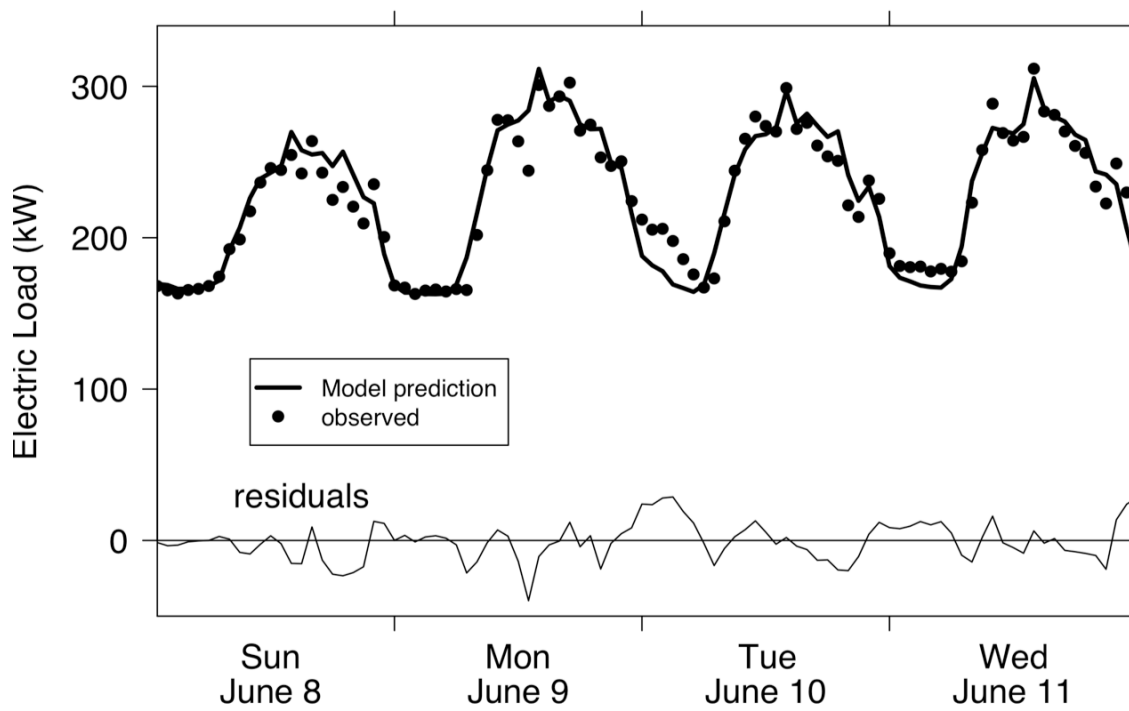


Figure 21: Predicted and observed hourly average load in Office Building A, for the same days as shown in the previous figure. Residuals are also shown.



Although some of the unpredicted variability in the load may be due to stochastic variation in the timing of the cooling load, it is obvious that there are other sources of variation between the model and the data. For instance, for several hours around midnight between Monday and Tuesday the load is about 30kW higher than predicted. Put another way, the “evening shoulder” is higher than expected on that Monday night.

As discussed above, while some of the features in Juvenile Hall repeat in the same 15-minute interval each day or each week, some loads are expected to vary in time. It may be essentially impossible to predict whether a building’s chillers are operating in a particular 15-minute interval on a warm day, while at the same time it may be certain that they will operate for approximately a certain fraction of a particular hour. If this is the case, then it may be possible to predict the hourly average load much more accurately than the 15-minute load. This can be done using the same approach as summarized in **Equation 5**, but using hourly averages of load and temperature rather than 15-minute data. The result for Juvenile Hall is shown in Figure 21 or the same four days in June 2008 that were discussed above. In this case, the residuals of the fit to hourly averages are only slightly smaller than the residuals of the fit to 15-minute data: for nine months of data, the standard error drops from 12.3 kW to 10.5 kW. This shows that in this building most of the error in fitting the 15-minute data comes from something other than slight time shifts of otherwise predictable loads.

As the above examples show, even the simple model described by **Equation 5** performs quite well for Juvenile Hall. However, there is one obvious model change that might be necessary to improve the model performance for buildings that (unlike Juvenile Hall) are unoccupied or sparsely occupied for part of each week. For such buildings, it makes sense to allow the temperature sensitivity (the coefficient $\beta_{T_{over}}$) to be different at different times during the day. This should improve the fit because, if the temperature setpoint changes during the day --- if the indoor temperature is allowed to be higher at night than during the day --- the temperature sensitivity should vary as well. To see why this is so, consider a case in which the interior of the building is held at 70F during the day, and is allowed to climb to 78F at night and on weekends. An outdoor temperature of, say, 72 F might lead to substantial daytime cooling requirements because of both the need to cool the ventilation air and because of internal loads, but the same temperature might lead to no cooling requirement at all at non-working hours (such as weekends or nighttime) because the internal loads are lower and because the setpoint temperature is higher.

We drop the low-temperature terms from the model for simplicity, which does not cause problems since, for the data that we are currently considering, there are few time periods when the outdoor temperature was below 50 F.

The model is easily modified to allow different temperature sensitivity during occupied and unoccupied periods. The term T_{over} defined following **Equation 5**, is used as a basis for defining two new temperature variables: $T_{Over,Occ}$ is equal to 0 when the building is unoccupied, and is equal to T_{over} at other times; and $T_{Under,Occ}$ is equal to 0 when the building is unoccupied, and is equal to T_{over} at other times. These temperature variables are used as follows:

$$\hat{y}(t_i, T(t_i)) = \beta_i + \beta_{T_{over,Occ}} T_{Over,Occ}(t_i) + \beta_{T_{over,Unocc}} T_{over,Unocc}(t_i) \quad \text{Equation 6}$$

Fitting this model with the data from Juvenile Hall reveals that the temperature sensitivity for unoccupied periods, $\beta_{T_{over,Unocc}}$ is estimated to be almost identical to the temperature sensitivity for occupied periods, $\beta_{T_{over,Occ}}$. This makes sense because this building is always occupied so thermostat settings are not changed in the building: there is no nighttime or weekend thermostat setback.

In contrast, consider Office Building A. The thin line in Figure 22 shows the observed load for four days in June 2008. There seems to be a morning catchup (rebound), at least on Monday and Tuesday, suggesting that there is a nighttime temperature setback in this building. The thick line shows the predicted load from **Equation 5**; this model assumes that the temperature sensitivity is the same for all time intervals, which is a poor assumption when thermostat settings are not constant. On Sunday, when the daytime temperature reached 90F, this model overpredicts the load. On Monday, when the daytime temperature reached 97 F, this model underpredicts the load.

Figure 23 compares observations to predictions from the model described by **Equation 6**, which estimates two temperature sensitivity coefficients, one for evenings and weekends and one for workdays. In this case, the sensitivity difference is quite substantial: during working hours, the building uses an additional 1.34 kW for each degree F that the outdoor temperature exceeds 65 F, whereas at other times it uses only 0.33 additional kW per degree F. The fit is improved overall: the standard error of the residuals is 9.1 kW, compared to 9.8 kW for the previous model. Furthermore, the residuals show less autocorrelation, that is, there are fewer cases in which several consecutive hours are over- or under-predicted.

[This space intentionally left blank so that Figures 22 and 23 can share the same page, next].

Figure 22: Load predictions and measurements for Office Building A, for 15-minute intervals, for several days in June 2008. Predictions are made using the model described by Equation 5 which assumes the same outdoor temperature sensitivity for all time periods.

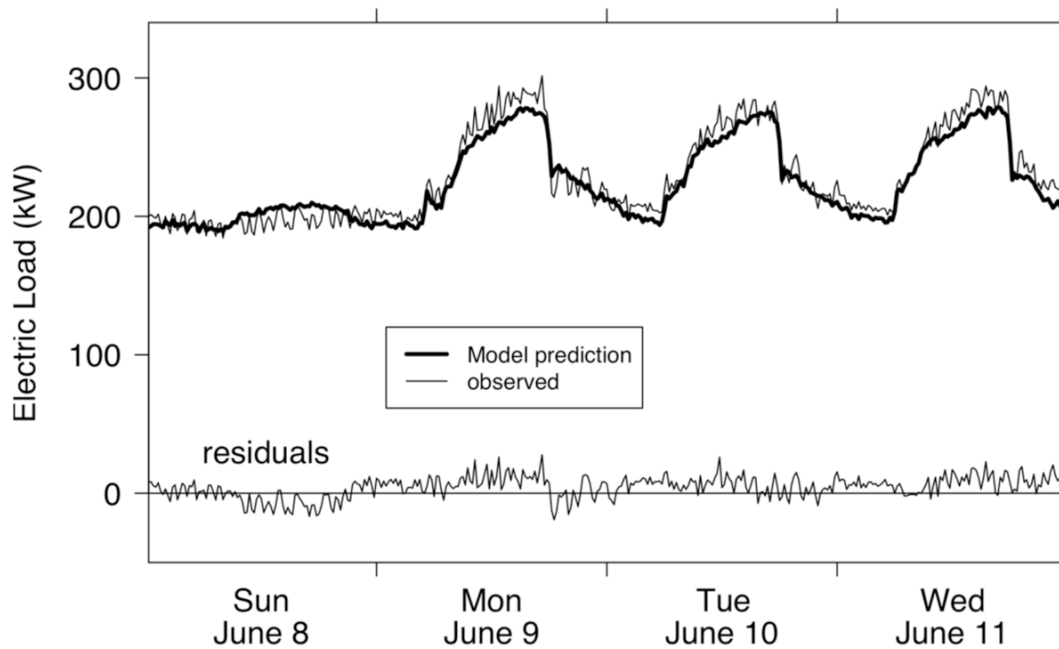
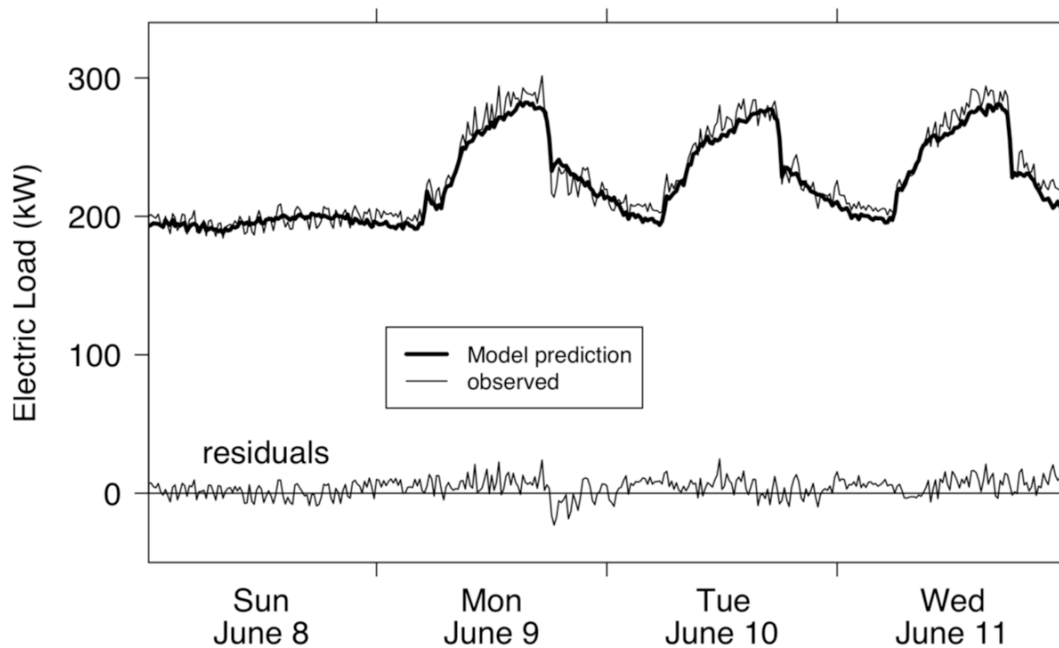


Figure 23: Load predictions and measurements for Office Building A, for 15-minute intervals, for several days in June 2008. Predictions are made using the model described by Equation 5 which estimates outdoor temperature sensitivity separately for time intervals when the building is and is not occupied.

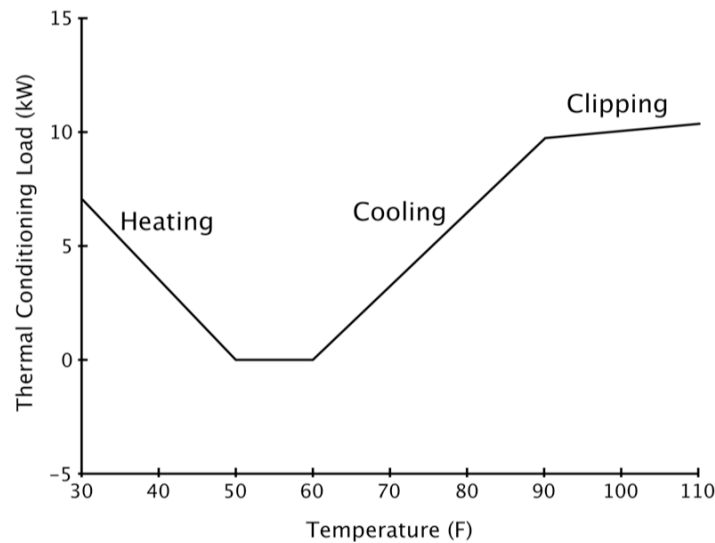


Given the large difference in temperature effect between working hours and other times, it might seem surprising that the model that uses only a single temperature sensitivity doesn't yield an even larger error than it does. The explanation is that the time-interval coefficients can partly compensate for the problems with the modeling of the temperature dependence. On Sundays, for example, the temperature-dependent term in **Equation 5** is always either too large (on hot days) or is correct (on cool days), so that *on average* the temperature-dependent term is too large on Sundays. The time-interval coefficients for Sunday are decreased to counteract this effect on average.

Even **Equation 6** doesn't perfectly capture the expected behavior of load as a function of temperature for times that the building has a high thermostat setting: the proper approach would be to define $T_{Over,Unocc}$ to reflect the higher setting; for instance, let $T_{Over,Unocc}=0$ for outdoor temperature less than, say 70 F (rather than for $T < 65$ F), and, for higher temperatures, to define it to be equal to the temperature minus 70 F rather than the temperature minus 65 F. In fact, even the reference values of 65 F or 70 F are just rough guesses of the outdoor temperature at which the building's cooling system has to start working. In principle, one could use a "change-point model" to estimate the reference temperatures, and this might be worth doing for certain purposes. However, the improvement from such a model will be small in most cases.

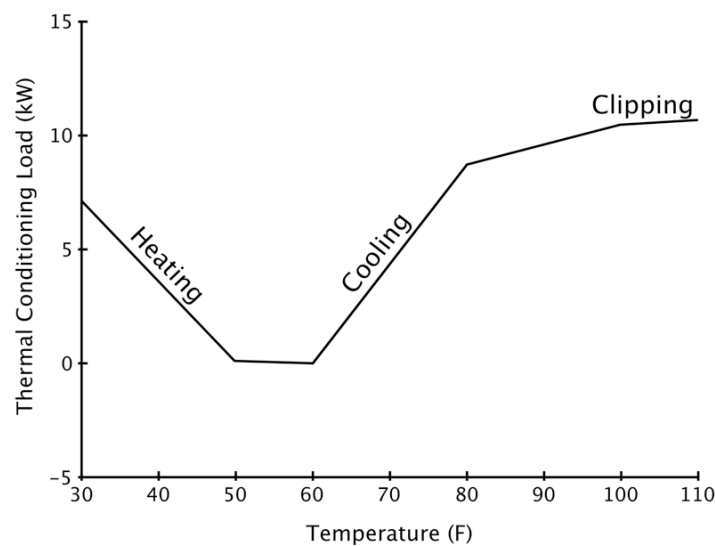
Simply allowing the effect of temperature on load to differ between working hours and other periods does not capture all of the systematic variation in temperature sensitivity that is expected. In fact, even at a given time of day the expected temperature dependence may be rather complicated. Figure 24 is a stylized plot of the expected relationship between load and outdoor temperature in a hypothetical building; at any given point, the slope of the relationship is what we call "temperature sensitivity." At low temperatures, the lower the outdoor temperature, the more energy the heating system consumes (e.g. for an electric heat pump); in this regime, electric load decreases with temperature. At cool temperatures, the building needs neither cooling nor heating. For moderate outdoor temperatures, the cooling load increases as the temperature increases, due to chiller operation, fan power, etc. In extremely hot conditions the cooling plant may be undersized, in which case its energy consumption will level out: once the system is working as hard as it can, its energy consumption cannot increase, a phenomenon we call "maxing out" or "clipping." (However, there could be additional loads, e.g. from the use of desk fans, that continue to increase with temperature).

Figure 24: Stylized plot of expected thermal conditioning load vs. outdoor temperature.



In fact, the situation is even more complicated than suggested in Figure 24. Many builders have multiple chillers. At low cooling loads, only one chiller operates. Chiller efficiency improves with load, so it is much better to have a single chiller operating at high capacity than to have multiple chillers at low capacity. If the temperature increases so that the first chiller is close to (but not at) its full capacity, a second chiller comes on, and both chillers operate at partial capacity. If the temperature continues to increase, eventually all of the chillers are operating at maximum capacity. Although chiller efficiency generally improves with load, some chillers are somewhat less efficient when working at full capacity than at some lower capacity. In any event, the fact that the efficiency varies with temperature suggests that a load model that is not linear in outdoor temperature might lead to improved fits. The situation is depicted in Figure 25.

Figure 25: Stylized plot of expected thermal conditioning load vs. outdoor temperature, showing the effect of decreasing chiller efficiency as the chillers approach full capacity.



The effects of maxing out and varying chiller efficiency are both minor from the standpoint of predicting long-term (weekly or monthly) building energy consumption: the effects are important only when the cooling system is operating close peak capacity, which is likely to happen only a few times per year because most cooling systems are sized to meet the demand on the hottest days.

However, although small on most days, the max-out and efficiency effects are not necessarily negligible, because the infrequent times at which these effects are largest are exactly those times when precise estimates are most needed to quantify the effectiveness of Demand Response, as we will discuss below.

A relationship such as the one depicted in Figure 25 can be fit using a linear model, by using the same sort of approach that lets separate slopes be used for heating and cooling, as described following **Equation 5**. Outdoor temperature is divided into bins, with bin j covering the outdoor temperature range from S_j to S_{j+1} . For simplicity of defining the rules, it is convenient if the lowest bin boundary is at 0 F and the highest is higher than the hottest temperature that is observed. A temperature variable T_j is defined from the outdoor temperature T as follows:

$$\begin{aligned} T_j &= 0 & T < S_j \\ T_j &= T - S_j & S_j \leq T \leq S_{j+1} \\ T_j &= S_{j+1} & T > S_{j+1} \end{aligned} \quad \text{Equation 7}$$

For instance, suppose there are bin boundaries at so $S_1=0\text{F}$, $S_2=50\text{F}$, and $S_3=60\text{F}$, $S_4=80\text{F}$, $S_5=100\text{F}$, and $S_6=120\text{F}$, and that the temperature at a given time is 73F. Then, by the rules given in **Equation 7**, $T_1=50\text{F}$, $T_2=10\text{F}$, and $T_3=13\text{F}$; $T_4=T_5=T_6=0$. Note that $T_1+T_2+T_3=73\text{F}$. Used together in a linear model, as in **Equation 8**, the T_j variables produce a piecewise-continuous model such as that illustrated in Figure 25.

$$\hat{y}(t_i, T(t_i)) = \beta_i + \sum_j \beta_j^T T_j(t_i) \quad \text{Equation 8}$$

The regression coefficient for bin j tells how much additional energy is used for each degree Fahrenheit within that bin. In the example in the previous paragraph, the coefficient for the first bin estimates the number of kW per degree for temperatures below 50F; the coefficient for the second bin estimates kW per degree for temperatures between 50 and 60 F; and so on. The temperature coefficients, taken together, characterize the sensitivity of the building's load to the outdoor temperature. Units are kW per degree. This way of characterizing temperature sensitivity differs from that of Coughlin et al. (2009): for each building, they ranked each day's electric load (e.g. day 1 had the 9th-highest load, day 2 had the 11th-highest load, ...) and ranked each day's peak temperature (e.g. day 1 was 14th-hottest, day 2 was 5th-hottest, ...) and correlated the rankings. A building in which the hottest days were also likely to be among the highest-load days was defined as temperature-sensitive, while a building in which the temperature rankings were not highly correlated with the load rankings was determined to be temperature-insensitive.

It is possible to add almost any number of additional variables, to create a more complicated (and realistic) model. We focus on four additional variables in particular. First, in addition to temperature, humidity also affects energy use, both because humid air has a different heat

capacity from dry air and because cooling air below its dewpoint temperature requires removing enough heat to cause the water to condense. Humidity is not usually a major factor in California, but it is not negligible. We include a single regression variable, $H(t) = T_{dew}(T) - T(t)$, where T_{dew} is the dewpoint temperature. When the relative humidity is high, this value is close to zero; when relative humidity is very low, the value is very negative. Empirically, this term allows the model to capture some of the effect of humidity on load.

Second, an examination of the residuals from any of the models discussed above reveals that there is autocorrelation of the residuals: if the predicted load at time t_i is higher than the measurement, then that is likely to be true at time t_{i+1} as well. In fact, the prediction for an entire day or week can be higher or lower than the observations. One way to adjust for this phenomenon is to include a different indicator variable for each week, so that just as each minute of the week has its own regression coefficient, each day of the year does as well. The week-of-year coefficients, β_{week} , quantify the amount by which a given week's average load is greater or less than predicted from all of the other variables. A summary of the magnitude of the week-of-year coefficients quantifies the amount of *unexplained* week-to-week load variation. It is possible to use this same approach to quantify unexplained day-to-day variation (or unexplained variation at other timescales).

Third, comparison of predicted to actual load reveals that the models discussed above do not predict enough morning catch-up (also known as “pull-down”) temperatures are high. The time-of-week coefficients that correspond to the catch-up period adjust for the average catch-up, so that, for instance, the power at 6:15 a.m. on Monday mornings is correct *on average*, but in fact the power requirement is much lower on cool mornings than on warm mornings. Rather than attempt to capture this phenomenon in detail, we implement a fairly simpleminded approach: we create yet another temperature variable, T_{morn} , that is equal to either the outdoor temperature minus 65 F, or zero, whichever is greater, for two hours after the building reverts to “occupied” mode, and is equal to 0 at all other times. The regression coefficient associated with this variable measures the additional energy used during this two-hour period only, per degree that the morning outdoor temperature exceeds 65 F.

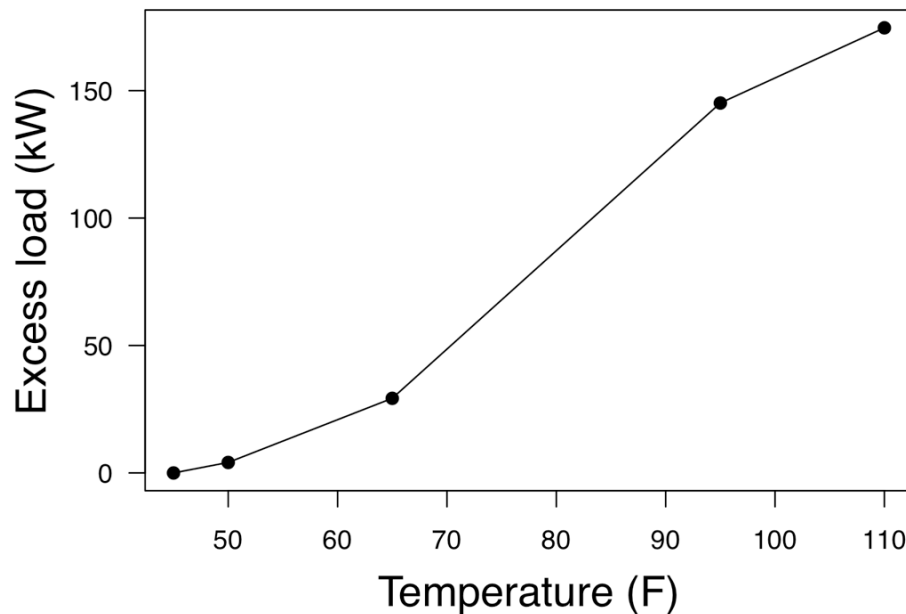
Fourth and finally, we need a way to estimate the effectiveness of Demand Response. DR is intended to decrease the load at specific times (such as between noon and 6:00 PM) on selected days, compared to what the load would be if no special actions were taken (see Motegi et al., 2007). To quantify the difference in load between DR days and typical days, we define two additional indicator variables, such that the regression coefficient for one of the variables quantifies the load reduction between noon and 3:00 PM, and the other quantifies the load reduction between 3 PM and 6 PM. The first of these, D_1 , takes the value 1 on DR days between noon and 3 PM, and 0 at all other times. The second, D_2 , is equal to 1 on DR days from 3 PM to 6 PM, and 0 at all other times.

$$\hat{y}(t_i, T(t_i)) = \beta_{week_i} + \beta_i + \sum_j \beta_j^T T_j(t_i) + \beta_H H + \beta_{morn} T_{morn} + \beta_{DR1} D_1 + \beta_{DR2} D_2 \quad \text{Equation 9}$$

The model described by **Equation 9**, includes (1) indicator variables for each day, to capture day-to-day variation that is not explained by other variables; (2) indicator variables for the time period during the week, to capture loads that repeat each week at the same time; (3) piecewise-linear terms to capture the effect of outdoor temperature; (4) a term that captures humidity effects that are linear in the difference between outdoor temperature and dewpoint temperature; (5) a temperature term that is nonzero only during the morning catch-up period (and is equal to the outdoor temperature during that period), to partially capture the temperature-dependence of the morning catch-up load; and (6) indicator variables that identify early- and late-afternoon DR periods, whose coefficients estimate the effect of DR.

The model described above was used to analyze data from an office building in Martinez, California. The resulting temperature-dependence of load is shown in Figure 10. We label the y-axis “excess load” rather than “cooling energy” because there is no certainty that all of the temperature-dependent load represents cooling energy: all we can say is that, statistically speaking, this is the component of load that varies with temperature. This building is discussed in more detail, below.

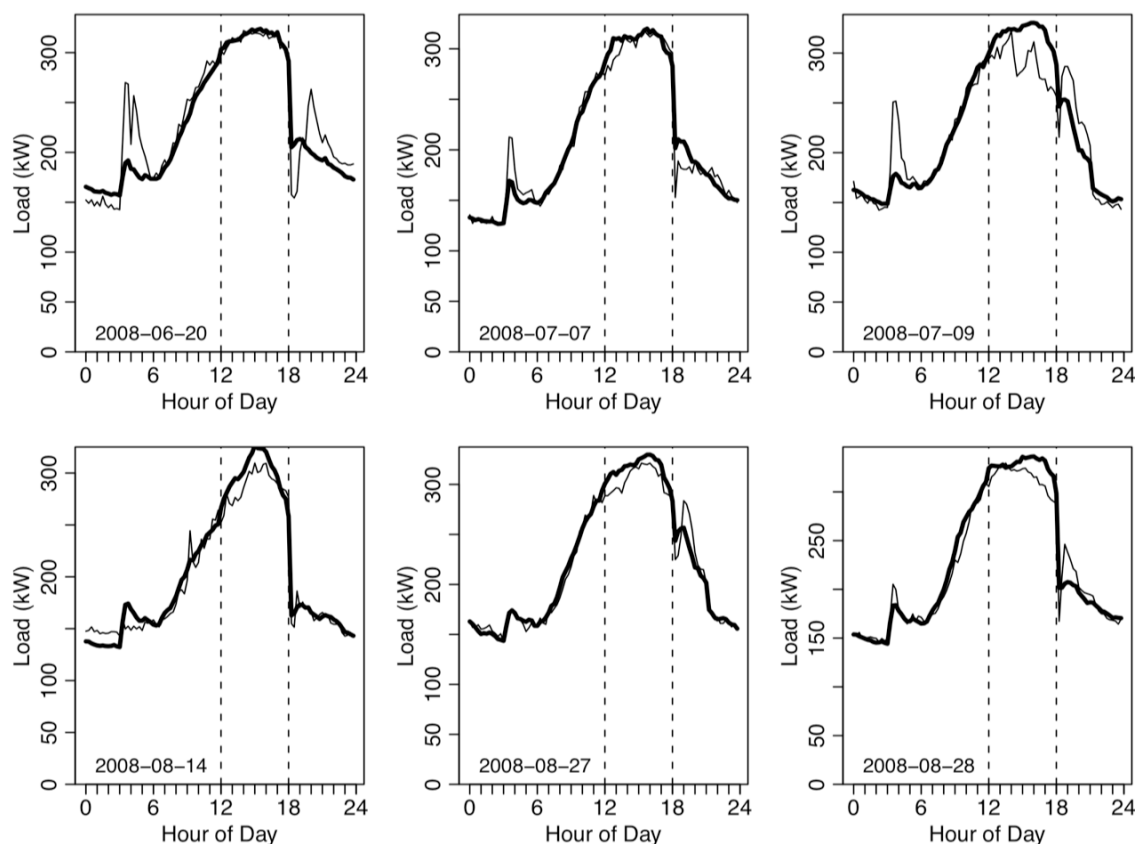
Figure 26: Excess load vs. temperature, for an office building in Martinez, California. Estimates are from a piecewise-linear model summarized in Equation 9.



Estimates of Demand Response effectiveness

Figure 27 shows model predictions and measurements from an office building in Martinez, California, for six DR days in summer, 2008. This is the same building whose temperature sensitivity is shown in Figure 26. The thick line shows the prediction *excluding* the DR coefficients, and the thin line shows the measurements. Since the DR coefficients are not included, the thick line shows the predicted load for a *non-DR* day that has the same explanatory variables (temperature, day of week, etc.) as the actual day. An effective demand response would yield an actual load less than the predicted load between 12:00 and 18:00. DR appears to have been ineffective (or perhaps was not even attempted) on June 20 and July 7, but was apparently modestly effective on the other days. Additionally, July 9, August 27, and August 28 all show an apparent rebound following termination of the DR event at 18:00.

Figure 27: Predicted (thick line) and observed (thin line) load in an office building on six DR days in summer 2008. In each case, the prediction ignores the effect of DR. The vertical dotted lines at 12:00 and 18:00 indicate the DR period.



The coefficients β_{DR1} and β_{DR2} directly estimate the amount by which load is decreased in the early afternoon and the late afternoon on a DR day. In this building, on an average DR day the load is decreased by 6.4 kW between 12:00 and 15:00, and by 7.5 kW between 15:00 and 18:00, compared to what would be expected for a non-DR day that is otherwise similar. (These represent about a 2% load reduction). The regression also reports a standard error of each of these estimates: for each period, the uncertainty is ± 1.2 kW. However, this uncertainty is an underestimate of the true uncertainty, because it assumes, incorrectly, that the model residuals have no autocorrelation. There are rather involved statistical methods that adjust for autocorrelation of the residuals and thereby give a correct uncertainty, but we do not apply them here; this might be a worthwhile area for future research.

Results from another building – this one just south of San Francisco Bay, in Milpitas California – are shown in Figure 28 and Figure 29. This building includes both manufacturing and office space. The building's DR performance is inconsistent: on July 9, August 28, and August 29 they achieved substantial load reduction; on June 20 and July 7 they appear to have had a slight load decrease; and on the other days it is not clear that they achieved any load reduction.

Figure 28: Predicted (thick line) and observed (thin line) load by hour of day in a manufacturing and office building, on eight days in 2008 during which Demand Response was performed from 12:00 to 18:00. The prediction does *not* include the effect of DR; successful DR would lead to load lower than predicted.

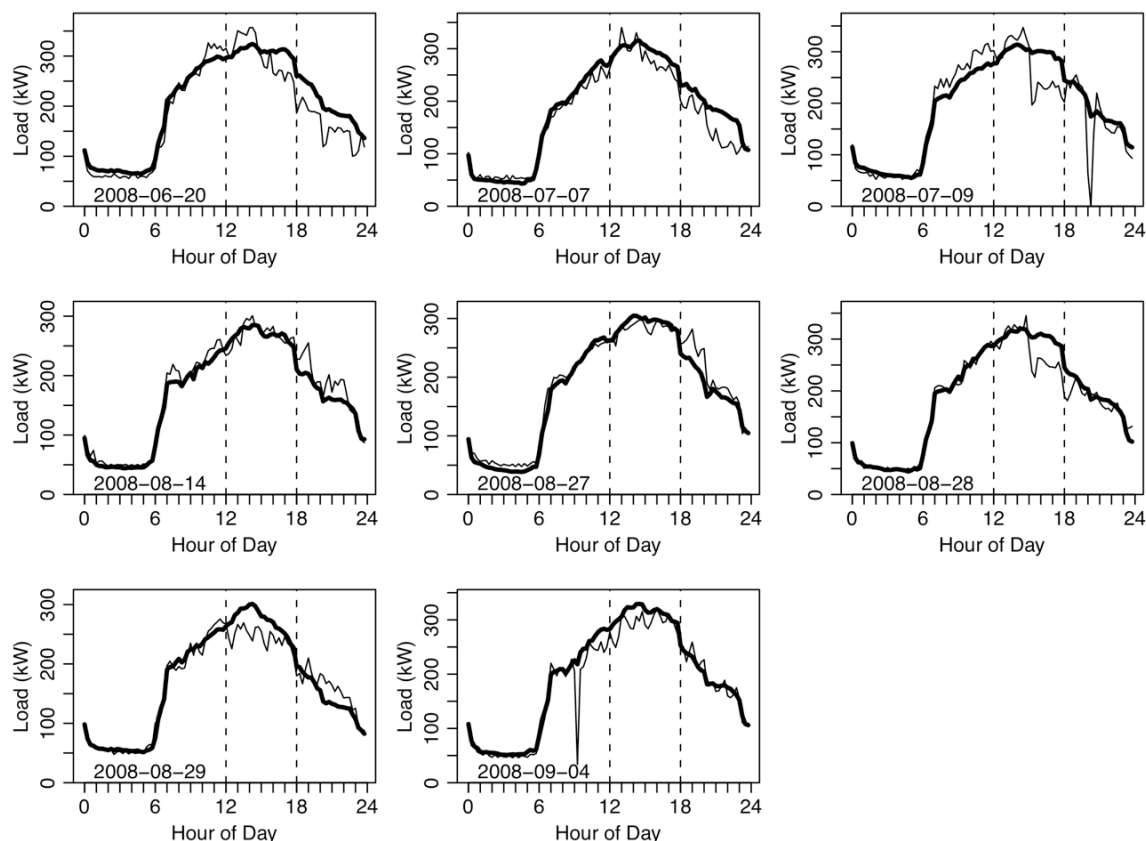
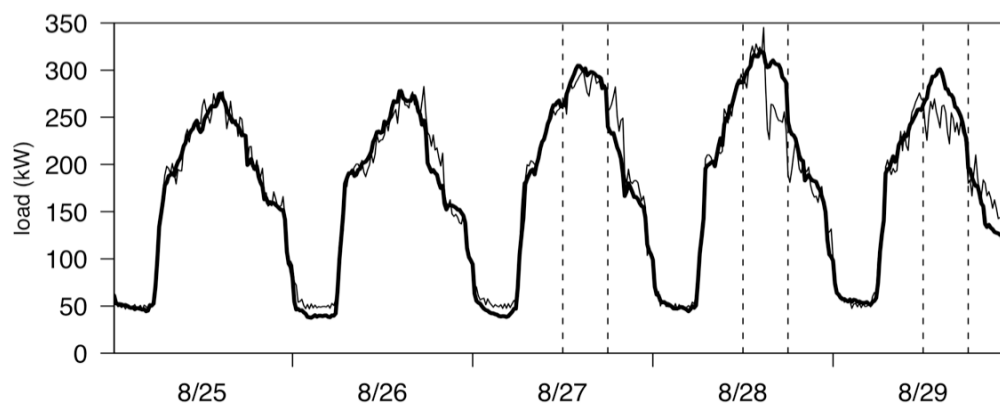


Figure 29: Predicted (thick line) and observed (thin line) load by hour of day in the same building that is considered in Figure 28 for a week in 2008 that included three DR days. The predictions do not include the effects of DR.



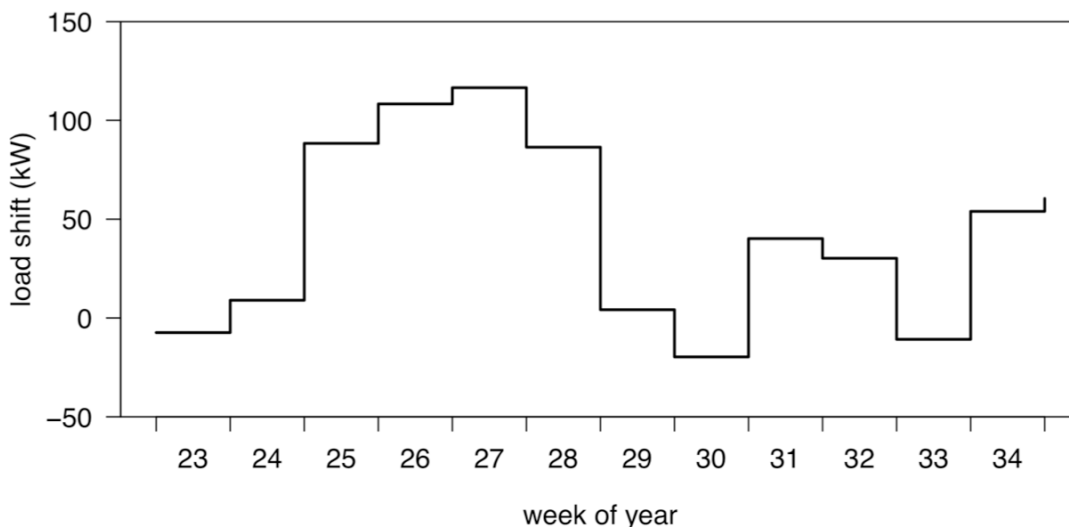
The model we have used in the examples above estimates the average DR effect, averaged over all of the DR events in the time period that is analyzed: that is the direct interpretation of the DR regression coefficients in **Equation 9**. It is also possible to make a separate estimate for each DR event, by creating a separate indicator variable for each event or by calculating the average residual during each DR period separately. The load reduction averaged over all events, as we have calculated above, can be estimated far more precisely than can the load reduction for any single event, so we recommend estimating the average DR effect rather than the effects for individual events unless the latter are needed for a specific purpose. Even estimating the average leads to large uncertainty (i.e. potentially large error): if the DR effect is small, of the order of a few percent reduction in load for example, then unexplained variation of approximately that magnitude (or larger) can cause a severe signal-to-noise problem.

Quantifying changes with time

There are several ways that changes with time can be identified; each has different advantages and disadvantages. Depending on the interests of the data analyst, it may make sense to use multiple approaches rather than rely on a single one.

A shift of weekly average load that is not explained by temperature or humidity is quantified by the β_{week} coefficients. For example, consider the load data from the Contra Costa Jail in June 2008 (see the lowest plot in Figure 10 and Figure 11). Although there are regular daily patterns in the load, there are substantial changes in load from one several-day period to the next that are clearly too large to be due to weather conditions or other available explanatory variables. The β_{week} coefficient is plotted, by week, in Figure 30. Recall that this coefficient estimates the magnitude of the load shift that is *not* attributable to changes in temperature or humidity. From the available data there is simply no way to determine the cause of this week-to-week variation, but at least its magnitude can be estimated. For this building the unexplained week-to-week variation is very large. Although in this case the week-to-week variation shows no particular pattern, this is not necessarily the case; for instance, a trend with time would show up as a trend in the weekly load shift.

Figure 30: Unexplained weekly load shift in the County Jail during the summer of 2008, from the model summarized in Equation 9.



An unexplained shift in weekly load is only one way in which the load shape can change with time. In fact, any of the model parameters can change: if the building's thermostat settings are changed, then the temperature sensitivity will change as well; if employees begin working longer or shorter hours, then the base load at a given time of the week will change; and so on. Although the week-to-week base load changes are automatically captured by the model, that is not true for the other parameters.

To quantify changes in model parameters, the model should be fit separately to different portions of the data --- e.g. for early and late summer, or separately for different years. Substantial changes in regression coefficients, such as coefficients of temperature, may indicate changes in the building's operation or equipment.

The regression models discussed above assume the times of day when the building is in "occupied" or "unoccupied" mode are known for each day. This is often the case in practice, because many buildings have a set schedule so that, for example, thermostat settings are changed at the same time each weekday. But of course the timing of these changes is under operator control and can be varied. Recognizing these changes is usually not difficult, as we will illustrate with real building data.

Figure 31: Electric load at 15-minute intervals for an office building in California, for the first full work week in June 2007 (top) and June 2008 (bottom). Monday is at the far left. Note the change in load on Saturday (between $x=5$ and $x=6$) from one year to the next.

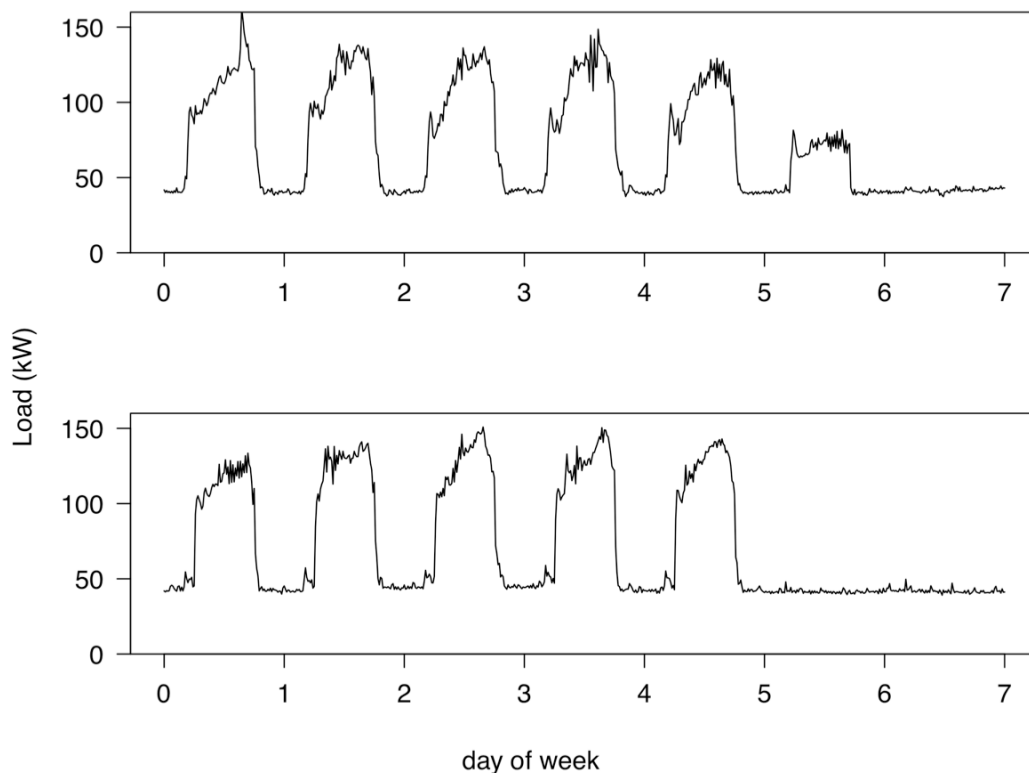
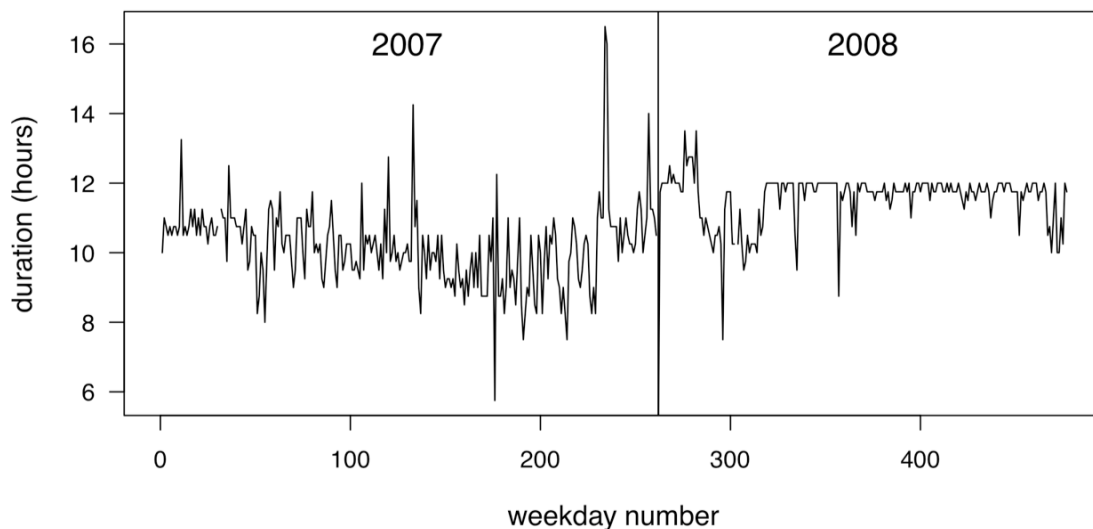


Figure 31 shows 15-minute electric load data from an office building in California. The upper plot shows data from the first full work week (starting on Monday) in June 2007; the lower plot shows data from the corresponding week in 2008. The base load, peak load, and overall load shape appear to be approximately the same for the two years. The most obvious difference is that the building uses the base load on Saturday in 2008, whereas in 2007 the load was elevated, although not as high as on the weekdays.

From Figure 31 alone, it would be easy to miss another substantial change from one year to the next: the building's hours with load substantially over the base were substantially extended in 2008 compared to 2007, and also exhibit much less day-to-day variation. Figure 32 shows the "high-load duration," defined above, for each weekday in 2007 and most of 2008. Most weekdays in 2008, the load was high for 12 hours, whereas a year earlier many weekdays had only about 9.5-10.5 hours of high load. In retrospect a change is visible even in Figure 31 – the building's load increases much more rapidly in 2008 than in 2007, so it more quickly reaches the point halfway to the peak – but noticing this fact, and figuring out when the change occurred, are hard to do from plots like Figure 31 and easy from Figure 32.

Figure 32: High-load duration in the office building of Figure 32Error! Reference source not found., for weekdays in 2007 and in most of 2008. In early 2008, the building's operating hours were extended by more than an hour each day, and became more uniform for different days of the week.

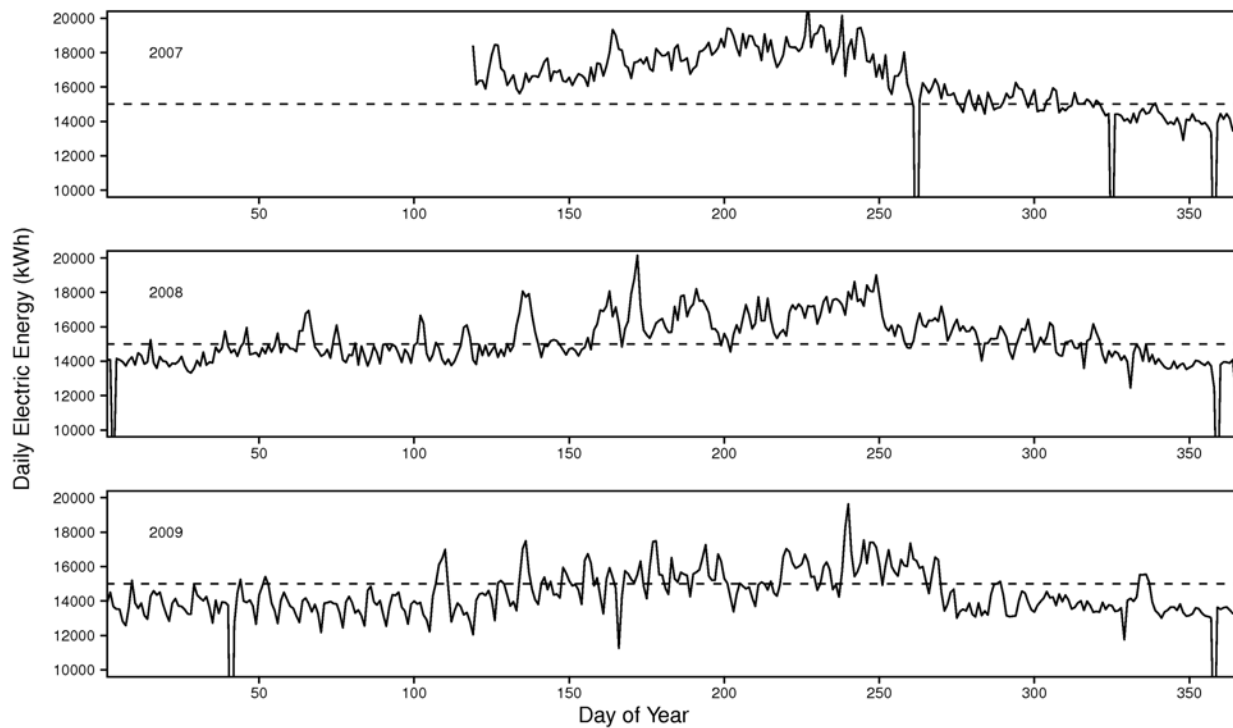


The salient change in load shape is that the steepness of the load curves: the load ramped up much more quickly in 2008 than in 2007. In 2007 the "high-load duration," as defined earlier in this document, often didn't begin until late in the morning ramp-up, after the morning catchup spike, whereas in 2008 the spike itself put the building into "high load." There are several possibilities for what might have changed in the building operation to cause this change in load shape, including a change in occupancy patterns (with more people arriving earlier in the morning), in lighting schedules, or in HVAC settings.

Example of analyzing year-to-year changes

We now illustrate briefly how the models and tools discussed in this report can be applied to quantify and investigate changes in electric load shape (i.e. load as a function of time). We use data from a furniture store (Ikea) in Palo Alto, California, from spring 2007 through the end of 2009. We begin by plotting daily electric energy (i.e. cumulative daily electric load), in Figure 33. As previously noted, such plots should use an aspect ratio such that key features are neither very close to vertical nor close to horizontal; for these plots, with 365 days of data, this means each year's plot should be short and wide.

Figure 33: Daily electric energy (kWh) at Ikea Palo Alto for each day from April 30, 2007 through December 31, 2009. A horizontal dashed line at 15000 kWh helps compare across years. Y-axis does not start at zero.



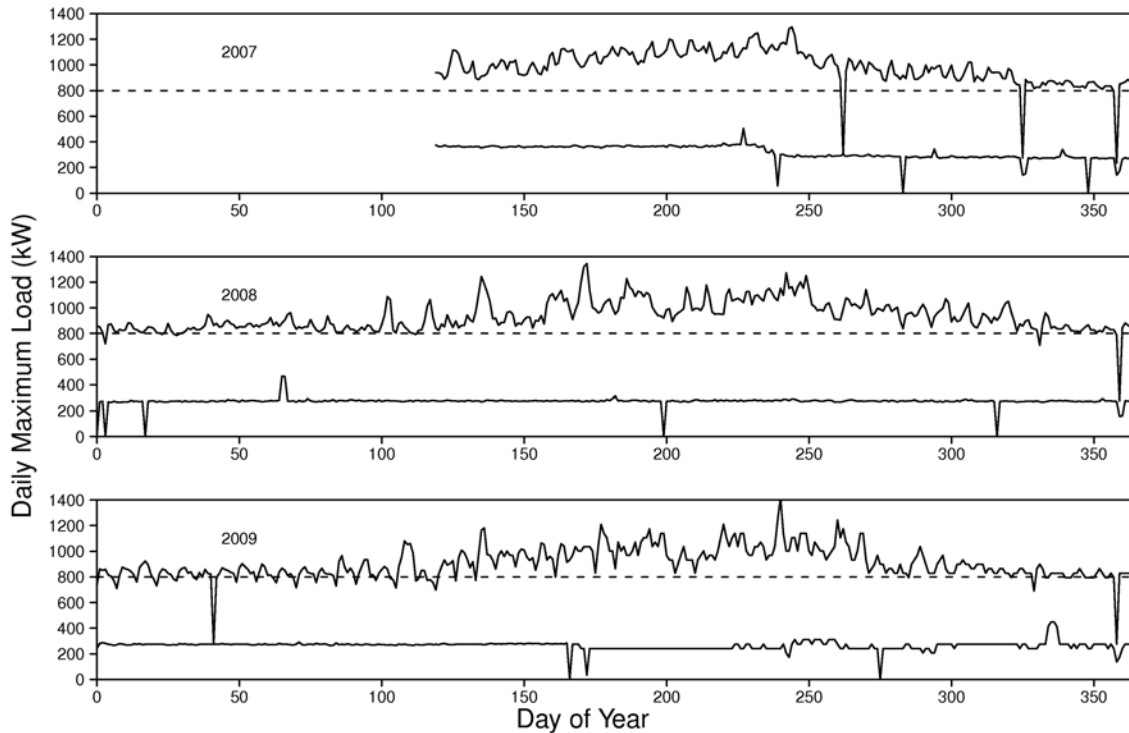
Energy usage dipped each year on December 25 (Christmas); in 2007 there are also dips on September 22 (which was Yom Kippur) and November 22 (Thanksgiving).

Each year, electricity consumption was higher during summer than during the rest of the year: each year there is a hump in consumption for days 150-275 or so. On average, consumption was lower in 2009 than in previous years as can be seen by noting the load curve compared to the dashed line at 15,000 kWh on each plot.

The next step is to look at base load and peak load by day, as in Figure 34. As discussed previously, for “peak” we actually use the 97.5th percentile, and for base we use the 2.5th percentile. In this case, we note two substantial shifts in base load: a large downward step around day 235 in 2007, and a smaller step around day 170 in 2009.

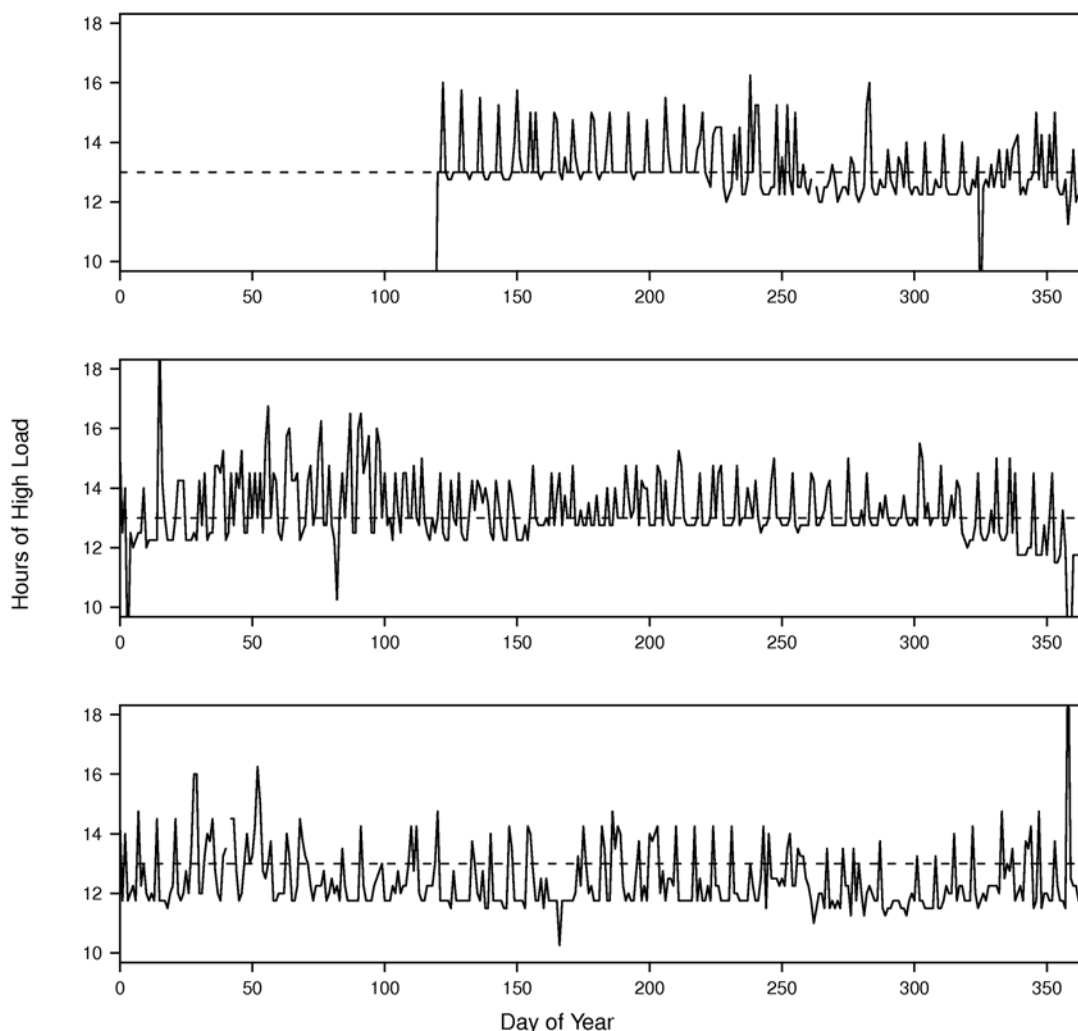
In contrast to the steadiness of the base load, peak load is rather variable. Especially in 2009, some of the variability is periodic, with a weekly peak and a weekly minimum; this is especially evident early in the year.

Figure 34: Daily maximum load and minimum load. A horizontal line at 800 kW helps compare load across years.



So far, we have looked at daily total electricity usage, and at the peak and minimum power consumed each day (i.e. the *magnitude* of power consumption); the next step is to look at *duration* of high power consumption. Figure 35 shows the number of hours of high load, as previously defined, for each day of each year. Although there is some irregularity, especially in early 2008, weekly variation is quite evident: one day each week has about 1.5 more hours of high load than the other days. Also, there were several substantial shifts in average behavior: (1) around day 250 of 2007, the number of high-load hours on most days decreased from 13 (indicated by the dashed horizontal line) to 12; (2) around day 50 of 2008, the hours stepped back up to 13; (3) around day 320 of the same year, the number of high-load hours stepped back down to 12 on most days, and remained there through 2009.

Figure 35: Hours of high load, by day of the year, for 2007 (top), 2008 (middle) and 2009 (bottom).



Using the simple graphical displays shown above, we see that:

1. Average daily energy use was lower in 2009 than 2007;
2. Base load was lower in 2009 than in 2007; and
3. The number of hours per day of high load was lower in 2009 than in 2007.

It's also interesting to look in detail at short (one- or two-week) periods from multiple years. Figure 36 shows the electric load vs. time for two weeks beginning with the first Sunday in June 2007, 2008, and 2009. Dashed lines at 400 and 1000 kW help compare one year to the next and one week to the next. At this level of detail, some day-to-day and week-to-week variation is evident (e.g. compare the first week in 2008 to the second). Other oddities are also apparent, such as the low load on the first Thursday in June 2009, and especially the unusual load shape on both Wednesdays in 2009 (which, in fact, occurs for other Wednesdays that summer as well).

Figure 36: Electric load vs time for two-week periods beginning with the first Sunday in June 2007, 2008, and 2009. Dashed lines at 400 and 1000 kW facilitate comparing one year to another.

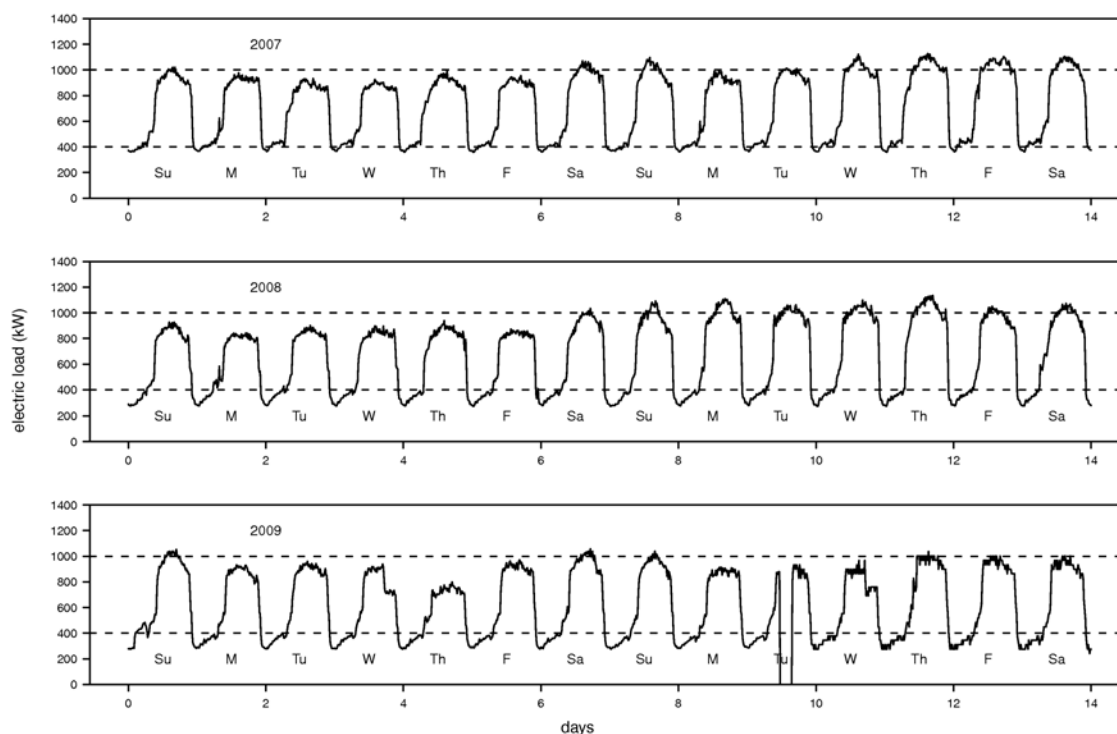
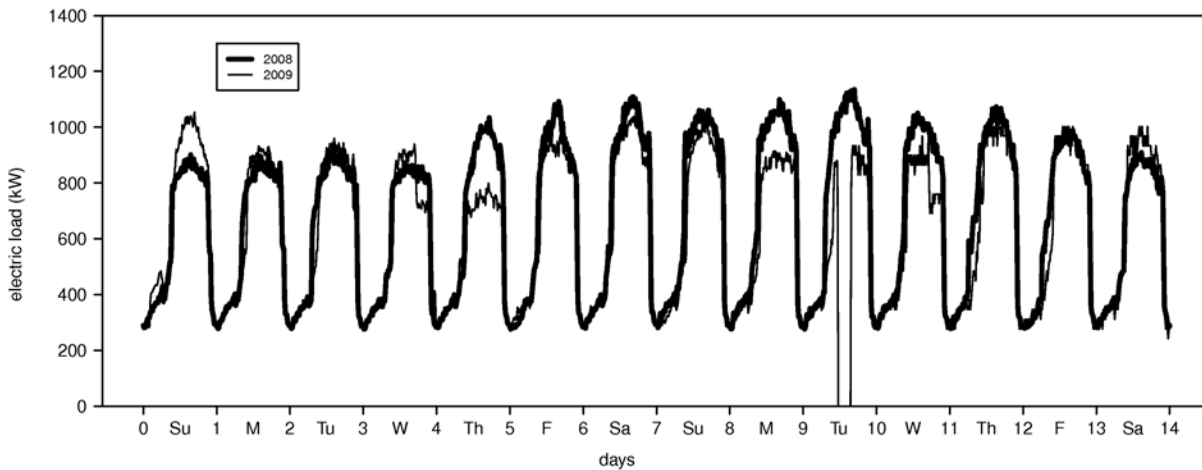


Figure 37 shows the same data as Figure 36, but only for 2008 and 2009, and with the years overlain on each other. At least for these two weeks, electric load was lower on most days in 2009 than in 2008 – the thin line is on, or below, the thick line almost everywhere. Although the load shapes in Figure 36 and Figure 37 are somewhat informative on their own, they are much more useful if analyzed in the context of a statistical model that adjusts for outdoor temperature: surely some of the variation from one day, one week, or one year to the next is due to different weather. In fact, for these weeks, temperatures were about 5C lower in 2009 than in 2008, and it is possible that this is a major contributor to the decreased energy use in 2009. So we will learn more by fitting a model to perform weather adjustment.

Figure 37: Electric load for the first two full weeks of June 2008 (thick line) and 2009 (thin line). Load is lower in 2009 than in 2008, but how much of this is weather-related?



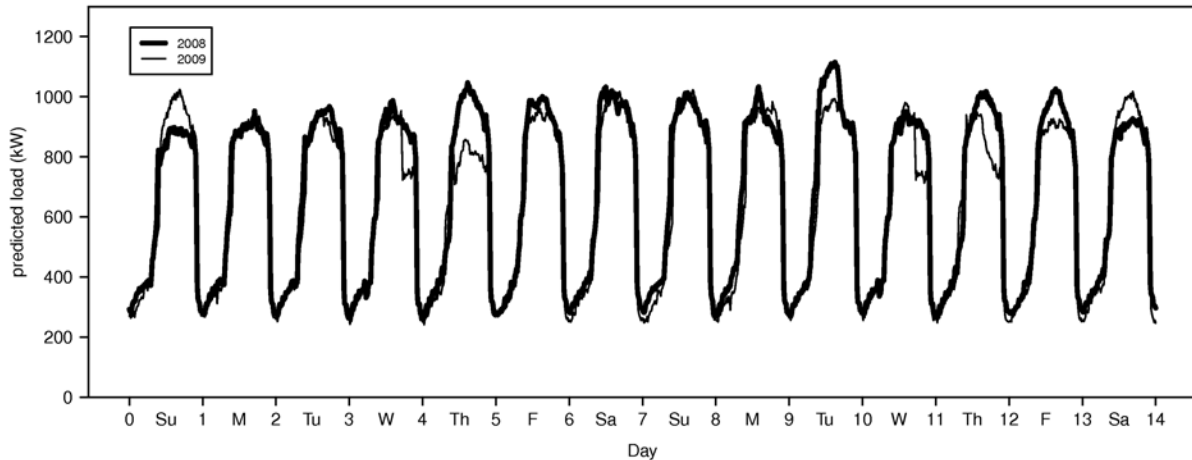
We fit the model described in **Equation 8** to data (including outdoor temperatures) from three-month periods in 2008 and 2009, to generate coefficient estimates associated with each 15-minute period of the week, as well as temperature coefficients for a piecewise temperature model. Week effects – the coefficients that adjust for otherwise-unexplained week-to-week variation in electricity usage – cannot easily be used in this case, because the week effects for 2008 would not, of course, be useful in 2009, so the models include only time of week, outdoor temperature, and humidity coefficients. These coefficients can be used to compare across years. For instance, we can use the model results from 2008 and 2009, together with actual temperature data from 2008, to predict how high the load would have been in each year if the temperatures had been the same in 2009 that they were in 2008. First, the 2008 load data and weather data (from three summer months, in this case) are used to estimate 2008 model coefficients, and these coefficients are then used with the weather data to predict the load. In essence, this just removes unexplained variation from the data, so that the predictions include only the systematic variation with time and temperature: otherwise we could simply use the data themselves rather than fitting the model. Next, the 2009 load data and weather data are used to estimate 2009 model coefficients, and these coefficients are then used with the 2008 *weather data* to predict what the load would have been in 2009, if the weather had been the same as in 2008.

Results for the first two full weeks in June are shown in Figure 38. We show only two weeks for illustration, but in fact the model encompasses the entire summer). Comparison to the actual data in Figure 37 reveals some interesting changes. On days 5-8 (Friday through Monday) the load in 2009 was noticeably lower than in 2008, but as Figure 38 shows this was largely due to weather differences between the two years. After temperature adjustment, the most prominent differences changes from 2008 to 2009 are a decrease in load on Wednesday afternoons, and most of the day on Thursdays.

Overall, though, it turns out that the impression of lower energy use in 2009 does reflect changes in building operation, not just milder weather: by using the coefficients obtained from the 2009 data, but substituting the 2008 weather, the predicted average load is 637W. That is, we estimate that if the building had its 2009 characteristics but had experienced 2008 weather, it would have averaged 637W. In fact, the building averaged 635W in 2009. That is, the difference between the 2008 weather and 2009 weather was small when averaged over the entire summer.

The 637 W predicted weather-adjusted average summer load in 2009 compares to the 2008 value of 674 W (both predicted and actual). So this store decreased its summertime energy use by about 5.5%, for reasons other than milder weather.

Figure 38: Model predictions for two weeks in June (starting with the first Sunday) for 2008 (thick line) and 2009 (thin line). In each case, the model coefficients are estimated using the load data and weather data from the year in question, but model predictions use the 2008 weather. Thus the 2009 predictions shown in this plot are the predictions for the energy use if the weather had been the same in 2009 as it was in 2008. Differences from Figure 37 are small but not negligible; compare the last two days on each plot, for example.



CHAPTER 4:

Conclusions

Whole-building electric load data can provide useful information about the operation of a building. Data from as short a period as a few weeks, can be used to estimate base load, peak load, duration of high load, and weather sensitivity, and these parameters alone may provide useful information to building operators and researchers.

In addition to results of statistical analyses, simple graphical displays that show load versus time, and facilitate comparison across days, weeks, months, or years, can also be useful. Data collected over a longer duration can be used to monitor changes in the building's operation, and to quantify the effectiveness of operational changes (such as demand-response) or equipment retrofits. In the text, we presented several rules of thumb for creating such displays.

Quantifying some parameters of interest, such as temperature sensitivity, requires fitting a statistical model. In this document, we have used ordinary linear regression, but with special features (time-of-week indicator variables, and piecewise-linear temperature dependence) that are more complicated than casual data analysts are used to. The level of complexity in fitting and interpreting statistical models probably limits the utility of these approaches to researchers, as opposed to building operators, at the moment. As discussed earlier in this report, additional work is needed to properly account for temporal autocorrelation so as to correctly assign uncertainties to parameter estimates. To convert regression models (or other statistical models) from research tools to tools that will be useful to building operators would require joint efforts of building energy researchers and software developers, to create a robust, easy-to-use interface for performing the analyses and summarizing the results in useful form. Also, ease of use would be greatly improved by automated selection of parameters that are currently the choice of the data analyst. For instance, for all of the buildings discussed in this report the data analyst first had to look at the load data to determine the starting and ending hours for the period during which the building is in its "occupied" mode, and then enter these hours into the data analysis functions; automating this procedure would improve convenience, but may not be trivial. For instance, hours of operation (or other building parameters) may change with time, perhaps more than once. Such changes can often be readily detected by a person through use of good graphical displays, as in Figure 35 for example, but automated recognition of these changes may not be so easy. Moreover, statistical models (and the tools to create them) should be developed that can handle the common case that some building parameters remain the same while others vary, rather than fitting a new, independent statistical model every time there is a change in a single parameter (such as an HVAC start time).

References

- Al-Hamadi HM and Soliman SA, Long-term / mid-term electric load forecasting based on short-term correlation and annual growth. *Electric Power Systems Research* 74:353-361, 2005.
- Cleveland, WS. *The elements of graphing data*. Hobart Press: New Jersey. 1994
- Coughlin K, Piette MA, Goldman C, Kiliccote S, Statistical analysis of baseline load models for non-residential buildings, *Energy and Buildings*, 41:374-381, 2009.
- Granderson, J, Piette MA, Ghatikar G, and Price PN. *Building Energy Information Systems: State of the technology and user case studies*. Lawrence Berkeley National Laboratory report LBNL-2899E, 2009.
- Pedersen L, Stang J, and Olsen R, Load prediction method for heat and electricity demand for the purpose of planning for mixed energy distribution systems. *Energy and Buildings* 40:1124-1134, 2008.
- Karatasou S, Santamouris M, and Geros V, Modeling and predicting building's energy use with artificial neural networks: Methods and results. *Energy and Buildings* 38:949-958, 2006.
- Kreider JF, Haberl JS, Predicting hourly building energy use: the great energy predictor shootout—overview and discussion of results, *ASHRAE Transactions* 100: 1104–1118, 1994.
- Motegi N, Piette MA, Watson DS, Kiliccote S, and Xu P. *Introduction to commercial building control strategies and techniques for Demand Response*. California Energy Commission, PIER. May 2007.